




National Ignition Facility Performance Review 1999



September 2000

Lawrence Livermore National Laboratory
University of California • Livermore, California • 94550

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This work was performed under the auspices of the U. S. Department of Energy by the University of California, Lawrence Livermore National Laboratory under Contract No. W-7405-Eng-48.

This report has been reproduced directly from the best available copy.

Available electronically at <http://www.doc.gov/bridge>

Available for a processing fee to U.S. Department of Energy
And its contractors in paper from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-mail: reports@adonis.osti.gov

Available for the sale to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-mail: orders@ntis.fedworld.gov
Online ordering: <http://www.ntis.gov/ordering.htm>

OR

Lawrence Livermore National Laboratory
Technical Information Department's Digital Library
<http://www.llnl.gov/tid/Library.html>



National Ignition Facility Performance Review 1999

John T. Hunt

September 2000

Table of Contents

1.0 NIF ENERGETICS UNDER UV DAMAGE AND BEAM QUALITY LIMITATIONS.....	1-1
1.0.1 Ultraviolet Component Lifetime Estimates	1-1
1.0.2 Recent Observations of UV Damage and Component Lifetimes	1-1
1.0.3 Discussion of the UV Damage/Lifetime Model and the Data That Goes into It	1-3
1.0.4 UV Lifetime Estimates and Parametric Study	1-6
1.0.5 Conclusion	1-11
1.1 Discussion of NIF Energetics Estimates and Safe Operating Limits.....	1-12
1.2 NIF Energetics Estimates for 11-5 Amplifier Configuration and Baseline Gains.....	1-15
1.3 Performance Risk Due to Low Amplifier Gain	1-22
1.4 Flowdown Analysis: 1 ω Laser Performance and Front-end Requirements.....	1-25
References.....	1-30
2.0 INJECTION LASER SYSTEM (ILS)	2-1
2.1 Introduction.....	2-1
2.2 MOR Overview	2-3
2.3 MOR Procurement Strategy	2-8
2.4 PAM Overview.....	2-9
2.5 MOR/PAM Prototype Integrated System Testing	2-11
2.6 PAM operation for ICF and non-ICF Missions	2-15
2.7 PAM 50 Shot Pointing and Far-Field Stability Experiment.....	2-17
2.8 PAM Procurement Strategy.....	2-21
2.9 PABTS System Overview	2-22
2.10 Procurement Strategy	2-26

3.0 AMPLIFIER	3-1
3.1 Primary Optical Performance Summary	3-1
3.2 Gain and Gain Profile	3-2
3.3 Prompt Pump-Induced Wavefront	3-7
3.4 Amplifier Thermal Recovery (Sutton).....	3-10
3.5 NIF Power Conditioning Status Summary	3-17
3.6 Amplab II and First-Article Tests	3-23
3.7 Flashlamp Damage to Laser Amplifier Slabs	3-25
 4.0 NIF FINAL OPTICS ASSEMBLY AND 3ω OPTICS	 4-1
Introduction.....	4-1
4.1 Frequency Conversion	4-3
4.2 Lifetime of Optics in Operation.....	4-14
4.3 Optics and Coatings	4-31
4.4 Summary—NIF Final Optics Assembly and 3 ω Optics	4-44
 5.0 WAVEFRONT CONTROL	 5-1
5.1 NIF Wavefront Control System.....	5-1
5.2 Deformable Mirror (DM).....	5-2
5.3 Wavefront Sensor	5-13
5.4 Wavefront Controller.....	5-18
 6.0 NIF OPTICS MANUFACTURING	 6-1
6.1 Properties of and Manufacturing Methods for NIF Laser Glasses	6-1
6.2 KDP Crystals.....	6-21
6.3 Developing Optics Finishing Technologies for the National Ignition Facility.....	6-38
6.4 NIF Optic Processing	6-59
6.5 Engineering High-Damage-Threshold NIF Polarizers and Mirrors	6-77

6.6 Diffractive Optics for the NIF.....	6-92
7.0 THE MISSION-FIRST DEPLOYMENT STRATEGY FOR THE NIF.....	7-1

1.0 NIF Energetics under UV Damage and Beam Quality Limitations

Clay Widmayer, Dave Munro, Steve Haan, John Hunt

1.0.1 Ultraviolet Component Lifetime Estimates

To keep NIF operating costs at a reasonable level, it is desirable to have ultraviolet (UV) optics last about 500 to 1000 shots between refurbishments. Experiments on Beamlet, SlabLab, and the Optical Sciences Laser (OSL) have indicated that it may be difficult to achieve the desired UV lifetimes with the expected NIF fluences of 8 J/cm^2 and above. Consequently, we have directed much effort towards measuring the UV damage properties of optical materials and developing theoretical models for predicting component lifetimes. In this section, we discuss our model for UV component lifetime and present some estimates for lifetime and operating costs that are relevant to the NIF.

The discussion is divided into three sections. In Section 1.0.2, we describe our recent experience with UV damage and explain why we think the damage will be a bigger issue with the NIF than it has been in the past. Section 1.0.3 includes a top-level discussion of the component lifetime model and the data that goes into it. Finally, in Section 1.0.4, we present some theoretical estimates for component lifetimes under conditions comparable to those on the NIF. We also include a parametric study estimating the sensitivity of UV lifetimes to variations in three key UV damage parameters: growth rate, growth threshold, and damage initiation. From the parametric study, we identify goals for improving the damage parameters that will allow us to achieve the desired UV lifetimes on the NIF.

1.0.2 Recent Observations of UV Damage and Component Lifetimes

Routine replacement and refurbishing of UV optics has always been a part of glass laser operation. However, our recent experience on Beamlet suggests that NIF conditions may produce types of UV damage that are qualitatively different from what we have seen in the past. Some of our historical perspective on UV damage is summarized in Table 1-1. The table contains information for systems ranging from Nova to our most recent UV experiments on Beamlet.

Table 1-1. Summary of historical data on UV damage.

Campaign	Date	# Shots	Typical Mean Beam Fluence (J/cm ²)	Comment
Nova	1985–99	14,000	1–2	Approx. 3 refurbishments
OMEGA	1995–	Not Available	1–2	
Beamlet 32 cm THG	8/94	30	6–8.7	FOA widely spaced in air Minor damage
Beamlet 37 cm THG	3/95	15	6–8.1	FOA widely spaced in air
Beamlet MDT	1/98	26	>5.7	FOA closely spaced in vacuum— Minor damage
Beamlet HDTII	7/98	23	>7.8	FOA closely spaced in vacuum— Major damage

As the table indicates, both the Nova and OMEGA lasers have been typically operated with UV fluences in the range of 1–2 J/cm², substantially lower than the highest mean fluences expected for the NIF. In the case of Nova, about 14,000 shots were performed in this fluence range, with only about three refurbishments for the frequency conversion crystals and the final focusing lens. The operating lifetimes seen on Nova would be considered more than adequate for the NIF. In fact, the operating cost model in the NIF conceptual design assumes a lifetime of about 1000 shots per optic.

The first system integrated tests using UV fluences comparable to those on the NIF occurred on Beamlet during 1994 and 1995. As shown in the table, two campaigns were conducted with a total of 45 shots being performed with mean fluences in the range 6–8.7 J/cm². A major difference between the configuration used in these tests and the NIF design was that the Beamlet UV optics were located in air as opposed to vacuum on the NIF. Also, the Beamlet optics were relatively widely spaced compared to the NIF Final Optics Assembly (FOA). Even with these important differences, the tests produced relatively minor damage while operating for the first time at UV fluences comparable to the NIF. Thus, the tests provided no clues that UV damage levels on the NIF would be qualitatively different than those we had seen in the past.

The first tests that really approximated the actual NIF operating conditions were conducted on Beamlet in 1998. In these campaigns, the UV optics were placed in vacuum in a close-packed configuration, similar to the current NIF design. Two series

of shots were conducted. In the first series, 26 shots were propagated with a maximum mean fluence of 5.7 J/cm^2 . In the second series, 23 shots with fluences up to 7.8 J/cm^2 occurred. In this series, major damage beyond what had been experienced on Nova or previous Beamlet experiments was observed. Therefore, the realization that NIF UV damage may differ qualitatively from past experience can be traced to the Beamlet experiments of 1998.

1.0.3 Discussion of the UV Damage/Lifetime Model and the Data That Goes into It

One of the notable results of the Beamlet experiments in vacuum was that many damage sites grew rapidly in size with each successive laser pulse. This observation alerted us to the fact that a realistic lifetime model must include two mechanisms for UV damage: damage initiation and damage growth. In the damage/lifetime models, initiation (number and initial size of damage sites) and growth parameters are experimentally measured quantities. These parameters are combined with a statistical description of the fluence distribution of a NIF beam. The NIF beam fluence distribution, compared to the optical damage parameters, is relatively well understood, both from Beamlet near-field measurements and from Prop92 two-dimensional fast Fourier transform (FFT) modeling. Likewise, damage initiation is a relatively well-studied topic. The biggest change in our view of UV damage is our understanding of growth. Consequently, much effort has gone into measuring this property in the past year.

Figure 1-1 shows UV damage growth rate measurements taken on OSL and SlabLab, the results of which are used in the damage/lifetime model. In the experiment, damage sites were initiated with a single, very high-fluence laser pulse of about 45 J/cm^2 , and the increase in size was measured due to subsequent irradiation at lower fluences. The figure on the left shows the increase in size with shot number for a damage site irradiated by an 8 J/cm^2 beam. The growth rate was fit with an exponential function. The figure on the right shows the measured exponential growth coefficient as a function of beam fluence. Experimentally, it is found that the area of a damage site after N shots at fluence F is given by:

$$A = A_i \exp\{2a(F - F_{th})N\} \quad (1)$$

where A_i is the initial size of the site, a is the exponential coefficient (equal to the slope of the data line in the right-hand figure), and F_{th} is the threshold for growth to occur (the x-intercept of the line). Experiments show that the coefficient, a , is between 0.03 and 0.05 per shot per J/cm^2 , and the threshold for growth is about 5 J/cm^2 . These two quantities, a and F_{th} , completely parameterize the UV damage growth.

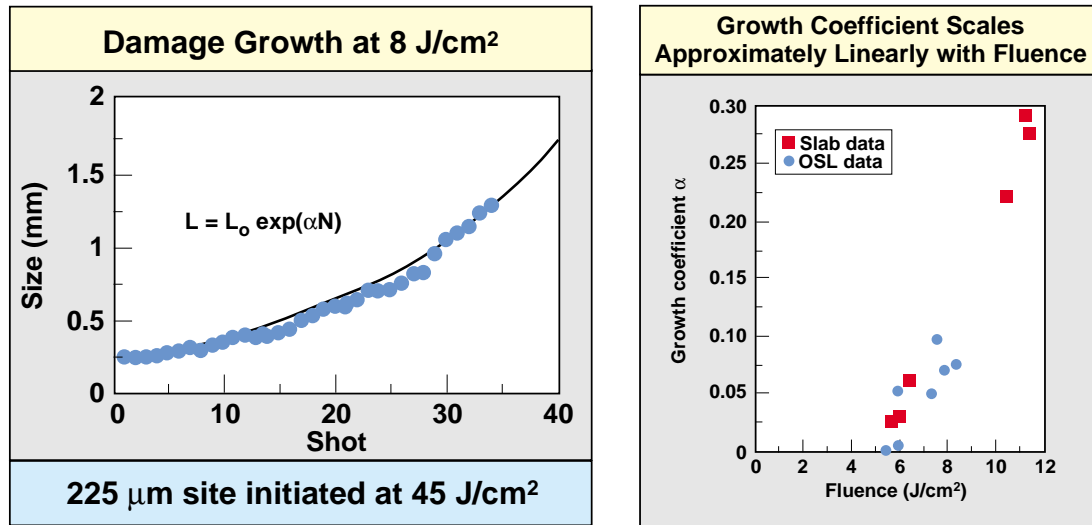


Figure 1-1. On the left, measured data for damage site diameter as a function of shot number at 8 J/cm². On the right, the measured exponential growth coefficient as a function of beam fluence. In the UV component lifetime model, damage growth is parameterized by the growth threshold, F_{th} , (the x-intercept on the graph) and the coefficient, a , (the slope).

An example of a damage initiation measurement is shown in Figure 1-2. The plot shows the number of damage sites per square centimeter (equivalently number per optic) as a function of irradiating fluence. The number of damage sites is a rapidly increasing function of fluence. We model the damage site concentration, or number per square centimeter, $C(F)$, with a function of the form:

$$C(F) = C_0(F/5)^7 \quad (2)$$




Figure 1-2. Measured data for UV damage initiation (number of damage sites per square cm) as a function of beam fluence. In the UV component lifetime model, we parameterize the initiation with the functional form $C(F) = C_0(F/5)^7$.

Using a power law for the fluence dependence of the concentration is equivalent to assuming the distribution of damage thresholds in the optic obeys Weibull statistics. Weibull statistics are believed to be a good description of the probability of rare event failure of mechanical systems and are widely used by engineers to simulate damage susceptibilities of materials. The Weibull exponent of 7 is found to give the best fit to the data in Figure 1-2. From the graph, a value of $C_0 = 0.004$ is used for the mean concentration at 5 J/cm^2 . In the parametric studies that follow, we parameterize the damage concentration by varying the prefactor, C_0 . It is notable that an improvement of as much as two orders of magnitude has been achieved in the damage concentration between 1998 and 1999.

Figure 1-3 shows an example of the statistical distribution used to model the fluence distribution of the laser beam. The picture on the left is a measurement of a Beamlet 3ω near-field profile. The picture on the right contains the fluence probability distribution for the same beam. Also shown in red in the picture on the right is a calculated probability distribution given by Rician statistics. The Rician distribution is an analytic form derived from the assumption that all modulation on the beam arises from random Gaussian phase errors, an assumption that is thought to be a good approximation for the NIF. As can be seen in the figure, the Rician distribution is a good approximation of the measured Beamlet statistics. The Rician distribution is completely parameterized by two quantities, the mean fluence and contrast ratio of the beam. This makes the distribution useful for parametric studies. The ability to simply parameterize the Rician beam is also useful since it provides an easy way to simulate increased beam modulations. This could be important if the NIF beam is more modulated than the Beamlet images predict. Factors that might cause the NIF's contrast to exceed that of the Beamlet measurements include increased noise in the NIF injection laser system and the limited resolution of the Beamlet diagnostic. In the studies that follow we typically model the NIF beam using a Rician distribution with a contrast ratio of 15%. The measured Beamlet beam, by comparison, has a contrast ratio of approximately 9%. Since the high-fluence tail of the distribution plays an important role in determining the UV damage and optic lifetime, the ability to treat the tails of the beam distribution realistically is important.

The damage initiation, damage growth measurements, and beam fluence statistics are combined to estimate the fraction of optical area damaged after N shots. We define the lifetime of an optic to be the number of shots until three percent of the beam area is obscured by damage. The fraction damaged, f_d , is obtained by performing a statistical integration over the beam fluence distribution, $P_B(F)$:

$$f_d = \int_0^{\infty} C(F) \pi r_0^2 \exp\{2a(F - F_{th})N\} P_B(F) dF \quad (3)$$

Here, r_0 is the initial radius of the damage site. Since $C(F)$ is the number of damage sites per unit area, πr_0^2 is the area of each, and $\exp\{2a(F - F_{th})N\}$ is the areal growth, their product is the fraction of area obscured by damage after N shots. The integral is a

statistical average over all the fluences in the beam. In practice, the lifetime is estimated by solving the above equation for N .

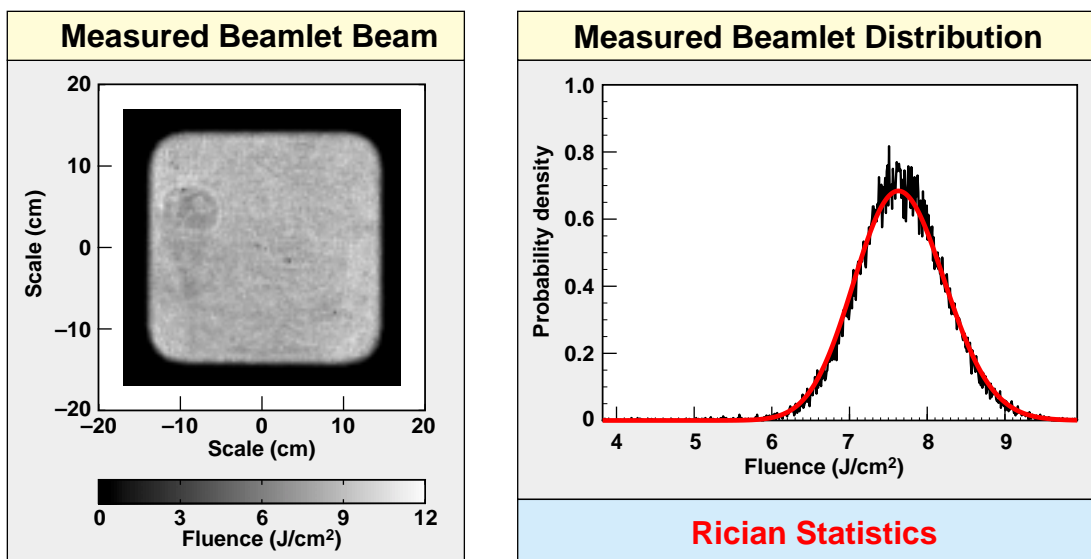


Figure 1-3. Example of a Rician fluence distribution used to model NIF beam statistics in the UV component lifetime model. The figure on the left shows a measured Beamlet near-field UV beam profile. The figure on the right shows the fluence distribution for the measured beam (black) and the Rician fluence distribution used in the model (red). The mean beam fluence and contrast ratio are variable parameters for the Rician distribution.

Note that the beam fluence distribution is the distribution for a single shot. The physical meaning of this is that the model simulates a situation in which the same beam profile is printed on the optic for N successive shots. Thus, all damage is initiated on the first shot, and only growth occurs on the subsequent $N - 1$ shots. This is consistent with what is observed on Beamlet. Also, in the model, the fluence fluctuations on the beam should be interpreted as remaining spatially fixed for all shots. That is, there is no randomization of the beam distribution from shot to shot. Again, this is not inconsistent with what we expect on the NIF since it is thought that the largest contributors to the NIF beam modulation will be surface finishing errors on optics, which are spatially fixed.

1.0.4 UV Lifetime Estimates and Parametric Study

An example of a UV component lifetime calculation is shown in Figure 1-4. In the solid curve, the damage and beam parameters were taken directly from the data in Figures 1-1 to 1-3. The growth coefficient, a , was taken to be 0.03 per shot per J/cm^2 , the growth threshold was $5.2 \text{ J}/\text{cm}^2$, and the laser beam fluence distribution was modeled with a Rician distribution with contrast ratio of 15%. The initial size of the damage sites was assumed to be $100 \mu\text{m}$ in diameter. The plot shows the calculated component

lifetime as a function of mean beam fluence. A second calculation (dashed line) is also included where the growth coefficient was reduced by 30X to show the effects of improving the growth parameters. Lifetime in these calculations is taken to be the number of shots until 3% of the optic is obscured by damage. Three percent was chosen because that is roughly the point at which optics are refurbished on Nova. The choice of 3% damage to define optics lifetime is somewhat arbitrary. However, because damage grows very rapidly during the last few shots, other reasonable definitions of lifetime, say between 1 and 5%, do not change the lifetime estimates appreciably.




Figure 1-4. UV component lifetime (number of shots until 3% damage) as a function of 192-beam energy. The calculation assumes UV damage parameters are given by 1999 measured values.

In the calculation using current growth estimates, the shot lifetime is less than 100 shots for all fluences greater than about 4.5 J/cm^2 or 1 MJ with 192 beams. For lower fluences, the lifetime improves rapidly as the beam energy is turned down, since very little of the beam is above the threshold for damage growth. Note that even at 4.5 J/cm^2 , where the mean of the beam is below the growth threshold, the lifetime is relatively short. Because of the hot spots in the beam, the mean beam fluence must be well below the growth threshold for long lifetimes to occur. In this case, the estimated lifetime is 1000 shots for a mean beam fluence of just less than 3 J/cm^2 . This is the fluence at which virtually all of the beam distribution is below the growth threshold of 5.2 J/cm^2 .

For fluences above 4.5 J/cm^2 , the component lifetime decreases steadily with increasing fluence. At 1.8 MJ (about 8 J/cm^2), the energy of the inertial confinement fusion (ICF) indirect-drive pulse, the lifetime is predicted to be about 20–30 shots. For weapons effects missions, even higher energies, up to 2.5 MJ, are requested, and the

model predicts lifetimes of about 10 shots. Of course, the NIF can still be operated with energies in the range 1.8–2.5 MJ. The lifetime estimates only indicate that shots in the high-energy range will be more expensive than our previous operating cost projections indicate if no further improvement is made in UV optics. Thus, we plan to increase the lifetime at high fluence by increasing the value of the growth threshold fluence and reducing the growth rate coefficient.

Figures 1-5 through 1-8 show the results of a parametric study estimating the sensitivity of UV lifetime to variations in damage parameters. In Figure 1-5 we plot the UV lifetime as a function of the growth coefficient, a . The three curves represent three different mean beam fluences; 4 J/cm² is in blue, 8 J/cm² in green, and 12 J/cm² in red. The green curve approximates the conditions of the ICF indirect-drive pulse, and the red curve roughly corresponds to high-energy weapons effects experiments. The current estimate of the growth coefficient, $a = 0.03$ per shot per J/cm², is at the far right of the horizontal axis. Lower values of a correspond to improvements in this parameter. As shown in the previous figure, with current estimates of growth, the model predicts between 100 and 200 shots at 4 J/cm² and 20–30 shots at 8 J/cm². Decreasing the growth coefficient by a factor of 30× to $a = 0.001$ would produce lifetimes of over 1000 shots at 8 J/cm² and several thousand at 4 J/cm². The very high-fluence weapons effects experiments with fluences of 12 J/cm² are predicted to have lifetimes of several hundred shots. Thus, a 30× reduction in growth coefficient would probably produce acceptable lifetimes for all operating fluences expected on the NIF.




Figure 1-5. UV component lifetime as a function of damage growth coefficient, a . Curves are shown for three different mean beam fluences: 4 J/cm² (blue), 8 J/cm² (green), and 12 J/cm² (red).

Figure 1-6 shows the change in UV lifetime due to improvements in the growth threshold, F_{th} . Again, lifetime is shown for three operating fluences: 4 J/cm² in blue, 8 J/cm² in green, and 12 J/cm² in red. The growth threshold is varied from its current

value of 5.2 J/cm^2 to a maximum value of 25 J/cm^2 . Lifetimes greater than 1000 shots are achievable for 4 J/cm^2 beams with an increase in growth threshold of about $1.5\times$ up to about 8 J/cm^2 . For 8 J/cm^2 beams, thousand-shot lifetimes can be achieved with an increase in F_{th} to 16 J/cm^2 or about threefold. With 12-J/cm^2 beams, the necessary increase in growth threshold is just less than fivefold to 24 J/cm^2 . In each case, the lifetime increases rapidly when the growth threshold is raised to a level about twice the mean fluence of the beam. This occurs when the majority of the beam fluence is below the threshold for growth. Thus, shot lifetime is very sensitive to changes in growth threshold.




Figure 1-6. UV component lifetime as a function of damage growth threshold, F_{th} . Curves are shown for three different mean beam fluences: 4 J/cm^2 (blue), 8 J/cm^2 (green), and 12 J/cm^2 (red).

In Figure 1-7, the shot lifetime is plotted as a function of damage initiation concentration. Three graphs are shown, one for each of the three beam fluences used in Figures 1-6 and 1-7. In each case, UV lifetime is plotted as a function of initiation in units of damage sites per optic. Our current estimates for damage initiation are about two sites with a 4-J/cm^2 beam, 150 with an 8-J/cm^2 beam, and 4000 with a 12-J/cm^2 beam. These estimates take into account the statistics of the beam and are based on a beam with a contrast ratio of 15%. The model predicts that lifetime is insensitive to variations in the number of damage sites. The weak dependence on initiation is a consequence of exponential growth; if growth is exponential, lifetime varies approximately logarithmically with initiation. This can be seen in the graphs in Figure 1-7, all of which show a slow increase in shot lifetime with reductions in damage concentration for site numbers greater than one per optic. For concentrations near or below one per optic, however, it is reasonable to hypothesize that defects can be reduced to zero by removing or etching out damage sites by hand. This is depicted in the graphs by sharp increases in the lifetime for concentrations below one per optic. If mitigating damage concentrations this way is feasible, then some relief from damage may be achievable through this method.




Figure 1-7. UV component lifetime as a function of damage initiation in units of damage sites per optic. Curves are shown for three different mean beam fluences: 4 J/cm² (blue), 8 J/cm² (green), and 12 J/cm² (red). In the calculations it is assumed that when defect concentrations are reduced to one site per optic, damage sites can be eradicated by hand to reduce the defect concentration to zero.

Finally, Figure 1-8 shows a calculation in which both the damage growth coefficient, a , and the damage growth threshold, F_{th} , were varied simultaneously. This calculation is probably more relevant to actual improvements in NIF optics than the previous cases, since we expect that future advancements in optics quality will improve all damage parameters. Naturally, improving growth threshold and growth coefficient simultaneously increases the shot lifetime fairly rapidly. Our goal of 1000-shot lifetimes at 8 J/cm^2 can be achieved with a simultaneous threefold improvement in growth and threshold.




Figure 1-8. UV component lifetime under simultaneous variations in both damage growth threshold, F_{th} and damage growth coefficient, a . Curves are shown for three different mean beam fluences: 4 J/cm^2 (blue), 8 J/cm^2 (green) and 12 J/cm^2 (red).

1.0.5 Conclusion

We have presented results based on a UV lifetime model incorporating damage initiation and growth that includes the effects of the beam fluence in a statistical way. Using 1999 estimates of growth and initiation, we predict lifetimes of a few hundred shots at 4 J/cm^2 , and a few tens of shots at 8 J/cm^2 . These are the lifetimes we would expect if no improvements are made in damage growth and initiation parameters. In Section 4, schemes for improving the lifetime of the optics are discussed. Using the model we can estimate what improvements are necessary to achieve our lifetime goals. Parametric studies suggest that either a thirty-fold reduction in the growth coefficient or a threefold increase in growth threshold would be sufficient to reach a thousand-shot lifetime at 8 J/cm^2 . The model predicts that a much weaker dependence of lifetime on damage initiation concentrations. However, if damage concentrations could be reduced to sufficiently low levels, perhaps a few per optic, damage sites might be effectively

eliminated by hand. The model predicts an improvement in concentration of two orders of magnitude would achieve this goal at 8 J/cm^2 .

1.1 Discussion of NIF Energetics Estimates and Safe Operating Limits

1.1.0 Introduction

In the previous section we discussed why the growth of optical damage in the UV section of NIF can limit the NIF's power and energy capabilities. Thus, during the first few years of NIF operation, it will be the growth of optical damage and the associated cost of refurbishing UV optics that will determine the number of high-energy shots that can be delivered to the target. Consequently, finding methods for reducing the damage growth rate and for increasing the damage growth threshold fluence have the highest priority in the NIF's technology development program. Although solutions will probably be of an interim nature with improvements coming over time, we believe that the damage growth problem will ultimately be solved.

Having said this, we now turn our attention to the energy and power limiters that will be encountered in the absence of UV damage growth. Here we discuss the NIF's power and energy capabilities for three types of pulse shapes (temporally flat, Gaussian, and ICF indirect-drive pulses) that are consistent with these limiters. These calculations include uncertainties related to our incomplete knowledge of amplifier gains and nonlinear effects such as self-focusing. This study is a follow-on to previous NIF performance work^{1,2} in which the following pulses were studied :

- A 1-ns, temporally flat pulse producing 750 TW relevant to Science Based Stockpile Stewardship missions.
- A 13-ns flat pulse producing about 2.5 MJ for Nuclear Weapons Effects Testing (NWET).
- A 1.8-MJ, 500-TW temporally shaped Haan pulse for ICF indirect-drive experiments.
- A 1.6-MJ, 440-TW pulse designed for ICF direct-drive experiments.

While these four pulses are thought to be representative of the NIF mission requirements, we expect a much wider range of power and energies will be required of the NIF. For example, it is expected that temporally flat and Gaussian pulses of various pulse durations, energies, and powers will be used on the NIF. Additionally, the ICF indirect-drive pulse is actually a family of pulse shapes, with continuously variable power and energy and shapes similar to the familiar 1.8-MJ, 500-TW Haan indirect-drive pulse.

The motivation and experimental basis for the various limits is discussed in Section 1.1. NIF performance estimates with the 11-5 amplifier slab configuration and baseline amplifier gains are presented in Section 1.2. In Section 1.3, we estimate performance risks due to gain uncertainties. Finally, in Section 1.4, the system's power and energy for the various pulses is flowed-down to two key system interfaces, the entrance to the frequency converter crystals (1ω output) and the main laser injection plane.

Four safe operating limits are suggested for the NIF. The limits are designed to limit “flaring” of the focal spot, prevent beam breakup in the 1ω and 3ω sections of the laser, and prevent overdriving the front end. Some of these effects, such as beam breakup and front-end drive, have been well studied. Therefore, we feel confident stating an operating limit for them. Other effects, such as spot size, are not as well understood. Operating limits based on these uncertainties we represent in the figures with bands.

1.1.1 ΣB -Integral Limit to Control NIF Focal Spot Size

To control the size of the NIF focal spot, it is desirable to specify an upper limit for the ΣB -integral accumulated in the main laser. The spot size depends strongly on ΣB because the “spot flaring” at high power is mostly due to nonlinear growth of low-frequency noise that passes through the NIF’s spatial filter pinholes. Hence, these noise sources undergo power growth, and a measure of that growth is due to the total B-integral (i.e., ΣB) accumulated from the main laser injection plane onward.

Two-dimensional FFT simulations of the NIF focal spot³ predict that the spot size increases rapidly for values of ΣB that exceed 4.5 radians. Hence, a prudent limit for ΣB is probably in the range of ~ 3.6 to 4 radians as measured between the injection plane and frequency converter crystals. With these values of ΣB , the FFT models predict that the 80% spot diameter is 300–350 microns and increases rapidly for larger values of ΣB .

Figures 1-9(a) and (b) show the ΣB -integral performance limit for temporally flat pulses. The green band represents those laser operating points with ΣB in the range 3.6–4 radians. Here, (a) shows the energy as a function of pulse length, and (b) shows power as a function of energy. The ΣB limit for Gaussian and Haan pulses is discussed in Section 1.2.




Figure 1-9. Comparison of NIF flat pulse operating limits for controlling focal spot size. The focal spot size limit is given by $\Sigma B \leq 3.6 - 4$ and is based on PROP92 modeling of the NIF. (a) The limits are displayed in UV energy pulse-length representation. (b) The information is the same as (a), but redrawn to show laser power as a function of energy.

1.1.2 ΔB -integral Limit to Control 1ω Beam Breakup

Beam breakup in the one-micron laser driven by self-focussing can result in damage to the one-micron optics. The parameter ΔB is a measure of nonlinear power growth and therefore, beam breakup. Tests on Beamlet and Nova⁴ revealed that a value of $\Delta B \sim 1.8$ radians roughly corresponds to the onset of severe beam breakup. As a result, $\Delta B \equiv 1.8$ radians is the limit we imposed in the following calculations.

1.1.3 3ω 750 TW Power Limit to Control 3ω Beam Breakup

The 3ω power limit is intended to prevent beam breakup and filamentation in the UV section of the laser. The value of 750 TW is based on filamentation experiments done on the OSL.⁵ The limit corresponds to a mean intensity of about 3.5 GW/cm^2 at the focus lens and was chosen produce an intensity-path length product (intensity times glass path) of less than 25 GW/cm .

1.1.4 2-J Injection Energy Limit

The 2-J limit is chosen to prevent over driving the front end and is based on prior NIF energetics modeling² that predict that 2 J is necessary to extract most of the energy from the main amplifiers. For higher output energies, injection energy requirements increase rapidly.

All of the NIF safe operating limits are shown in Figures 1-10(a) and (b). The 1ω ΔB -integral limit is shown in red, the 750-TW limit is in blue, and the 2-J injection limit is in yellow. The laser performance within these limits is discussed for flat, Gaussian, and Haan pulses in Section 1.2.




Figure 1-10. Comparison of all suggested NIF flat-pulse operating limits. The limits are for controlling 3ω beam breakup (UV power less than 750 TW, shown in blue), 1ω beam breakup ($\Delta B \leq 1.8$, shown in red), injection energy less than 2 J (shown in yellow), focal spot size ($\Sigma B \leq 3.6 - 4$, shown in green). (a) The limits are displayed in UV energy pulse-length representation. (b) The information is the same as in (a) but redrawn to show laser power as a function of energy.

1.2 NIF Energetics Estimates for 11-5 Amplifier Configuration and Baseline Gains

1.2.1 Temporally Flat Pulses

A complete picture of the NIF flat-pulse operating limits discussed in the previous section is shown in Figures 1-11(a) and (b). The plots show the safe operating limit obtained by combining the individual limits. The combined curve is obtained by taking the lowest value of all of the individual curves. The sections of the operating curve are color coded to indicate which effect limits performance in that region: black for UV damage initiation, red and green for ΣB and ΔB respectively, blue for 3ω beam breakup and yellow for injection energy. Again, the performance is depicted in the form of energy vs pulse length in Figure 1-11(a) and power vs energy in 1-11(b). In each case, a band is shown to depict the range of uncertainty in the ΣB operating limits.

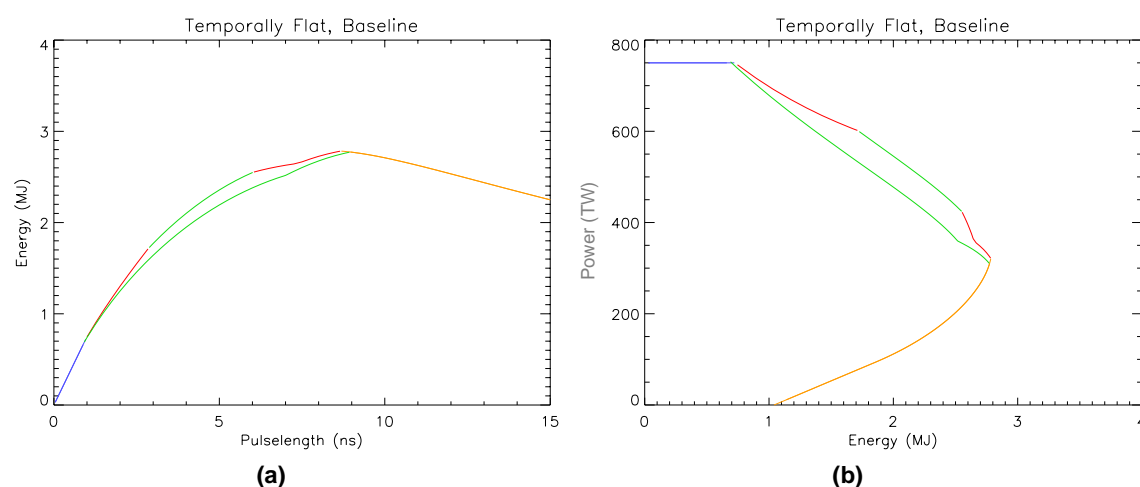


Figure 1-11. Suggested overall NIF flat-pulse operating limit. The limit is obtained by combining the four individual limits in Figure 1-10. (a) The limit is shown in UV energy pulse-length representation. The width of the operating band is due to uncertainties in the focal spot size limit. (b) The information is the same as in (a) but is redrawn to show laser power as a function of energy.

The maximum flat-pulse energy consistent with the operating limits is about 2.7 MJ, a value that is obtained at a pulse length between nine and ten nanoseconds. For longer pulselengths, the 2-J injection limit decreases because increasing the pulse length decreases the conversion efficiency, thus reducing the amount of 3ω energy obtainable from 2-J input. Therefore, once the maximum energy is reached, increasing the pulse length decreases the energy on target due to reduced conversion efficiency.

An important property of the temporally flat pulse is its high maximum energy, 2.7 MJ. Because the flat pulse is optimal for frequency conversion, more energy is obtainable with it than with other pulse shapes. The maximum energy occurs, however, at a relatively low power: about 280 TW. In the next section, it is shown that the Gaussian pulse achieves its maximum energy at a higher peak power but lower energy than a flat pulse. This property might make the Gaussian pulse useful for applications in which high power and energy are desired simultaneously.

Figure 1-12 compares the NIF flat pulse performance to requests of the Stockpile Stewardship High Energy Density Experimental Science (HEDES) and Nuclear Weapons Effects Testing (NWET) user communities. In general, experiments by both these communities will be more effective given more power and/or energy, and so the experiments will be designed at the limits of the NIF capabilities. As shown in Figure 1-12, different categories of experiments will operate at various limits of the power/energy curve:

For high-temperature hohlraum experiments, maximum 3ω power is the overriding requirement. Given maximum power, the experiments will probably be performed at the maximum energy consistent with that power. The 3ω beam breakup limit sets the maximum power, and either a 1ω B-integral limit or a pessimistic UV damage initiation limit sets the maximum energy that can be delivered at that power.

An explicit UV damage initiation limit is not shown in Figure 1-12. However, the range of experiments in the figure was derived from a damage initiation limit of $8\text{--}9\text{ J/cm}^2$ with a 3-ns Gaussian pulse and square root of pulselength scaling for other pulses. As discussed in Section 1.0, we have since learned that UV damage growth, not initiation, is more important for determining UV component lifetime and operating costs. Consequently, the $8\text{--}9\text{ J/cm}^2$ damage initiation limit is probably not as relevant as we had previously thought. So, when UV damage growth is brought under control, system power and energy will probably be determined by a B-integral limit as shown in Figure 1-12. However, the experimental operating points shown in the figure (and also Figure 1-15) date from the earlier assumption of the UV damage initiation limit.




Figure 1-12. Comparison of the overall NIF flat-pulse operating limit with the flat-pulse requests of the NIF user community. The operating limit and user requests are shown in UV power-energy representation. The user requests are expected to require operation at a large variety of power-energy points along the NIF operating limit.

Radiation-flow and hydrodynamics HEDES experiments will be performed at the highest possible energy, since these experiments can best be done and diagnosed at high energy. In this case, the extreme operating point is set by the UV damage initiation and the injected energy limits. Any reduction in the UV damage operating limit will have a direct impact on the quality of these types of experiments. Note that the injected energy limit is set by the low-intensity conversion efficiency and might be improved by adding a conversion crystal.

Indirect-drive equation of state (EOS) experiments will be distributed along the UV damage initiation operating limit. The quality of these experiments will be directly affected by reduction in that limit. Direct-drive EOS experiments are least affected by NIF operating limits, although the end-point of these experiments will be determined by the UV damage limit.

The NWET x-ray source experiments are constrained by NIF operating limits in much the same way as the HEDES indirect-drive EOS experiments: the UV damage initiation will set the operating limit. (A second category of NWET experiments, not shown, utilizes ignition and has the same requirements as ignition.)

1.2.2 Gaussian Pulses

Figures 1-13(a) and (b) show the individual NIF performance limits for Gaussian pulses. The limits are defined the same way as for temporally flat pulses. The motivations for the limits were discussed in Sections 1.1.1 to 1.1.5. However, because certain laser parameters such as UV damage, B-integral, and frequency conversion depend on the details of the pulse shape, the power and energy performance of the system is different for Gaussian and flat pulses.

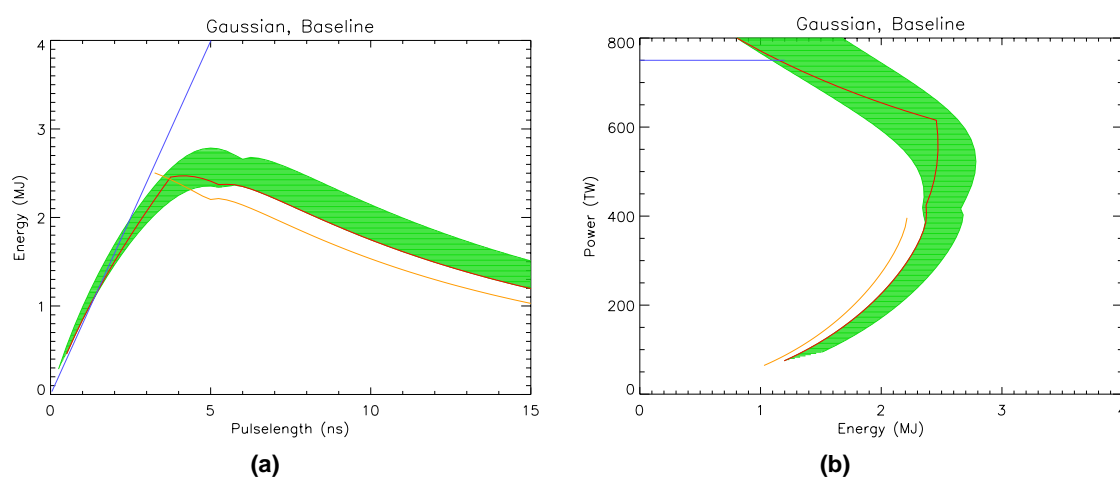


Figure 1-13. Comparison of all suggested NIF Gaussian-pulse operating limits. The limits are for controlling 3ω beam breakup (UV power less than 750 TW, shown in blue), 1ω beam breakup ($\Delta B \leq 1.8$, shown in red), injection energy less than 2 J (shown in yellow), focal spot size ($\Sigma B \leq 3.6 - 4$, shown in green). (a) The limits are displayed in UV energy pulse-length representation. (b) The information is the same as in (a) but redrawn to show laser power as a function of energy.

The combined operating limits are shown in Figures 1-14(a) and (b). The Gaussian pulse provides more energy at moderate pulse lengths (three to five nanoseconds) than the flat pulse. The increased energy output for this range of pulse lengths is a result of the fact that UV damage and B-integral effects are more favorable for the Gaussian pulse than for the flat pulse. This effect is a trade-off against the reduced power and energy obtainable from the Gaussian pulse at long pulse length (low power). The reduced performance at long pulse length is a consequence of the fact that Gaussian pulses have less desirable frequency conversion properties than flat pulses. In summary, the Gaussian pulse can be thought of as producing higher power at pulse lengths between 3 and 5 ns, while the flat pulse produces more energy at long pulse lengths.

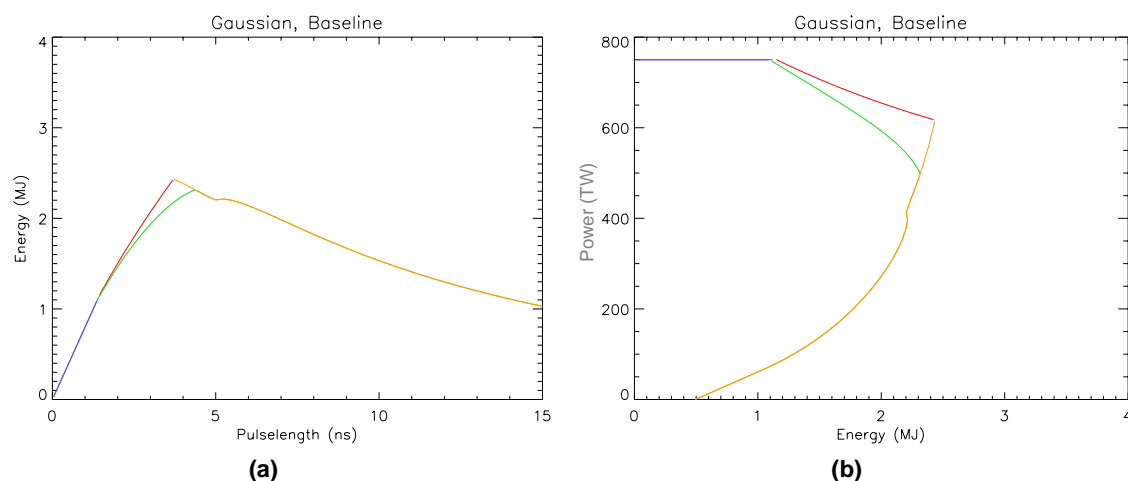


Figure 1-14. Suggested overall NIF Gaussian-pulse operating limit. The limit is obtained by combining the four individual limits in Figure 1-14. (a) The limit is shown in UV energy pulse-length representation. The width of the operating band is due to uncertainties in the focal spot size limit. (b) The information is the same as in (a) but redrawn to show laser power as a function of energy.

Figure 1-15 shows the Gaussian laser performance curve. As reference, the HEDES and NWET user requests are shown carried over unchanged from the square-pulse Figure 1-12. The comparison helps make it evident which types of experiments might be improved by the use of Gaussian pulses.




Figure 1-15. Comparison of the overall NIF Gaussian-pulse operating limit with the weapons effects testing requests of the NIF user community. The operating limit and user requests are shown in UV power-energy representation.

For some cases, Gaussian pulses will be inferior to square pulses. For other reasons characteristic of the experiment design, high-T hohlraums and direct-drive EOS probably require flat pulses. The laser constraints for Gaussian pulses are, hence, likely irrelevant to these experiments. For radiation flow and hydrodynamics experiments, flat pulses will probably be superior, because the requirement for maximum possible energy is better met with square pulses.

On the other hand, for indirect-drive EOS and NWET x-ray source experiments, Gaussian pulses may give superior experiments. In these cases, the power and energy could be increased beyond the area indicated in Figure 1-15, up to the limit imposed by damage initiation. (Of course the operating areas would then be moved up from those indicated in Figure 1-15, which are carried over directly from the square-pulse figure.) Note that the pulse shape does not actually need to be Gaussian to derive this benefit: any pulse-shape with power that decreases in time, so that a significant fraction of the energy is delivered at much less than peak power, will show a similar effect. For hohlraum targets like indirect-drive EOS, such a decreasing pulse can efficiently hold the hohlraum at constant temperature, unlike a flat pulse that results in a slowly rising hohlraum temperature. The NWET x-ray source targets may also be able to use decreasing pulses.

1.2.3 ICF Indirect-Drive Haan Pulses

Figures 1-16 and 1-17 show the laser performance calculated for ICF indirect-drive Haan pulses. The individual laser safe operating curves are shown in Figure 1-16. The combined curve is shown in Figure 1-17. The upper edge of the performance band in

Figure 1-17 represents the most optimistic values for UV damage initiation and Σ -B limits, and the lower edge of the band represents the most pessimistic. Also shown on the graph are contours (“bird’s beak” curves) that depict regions of constant ignition probability. For reference, the 1.8-MJ, 500-TW baseline Haan pulse is shown by an asterisk.

Each ignition probability contour consists of two branches (forming the upper and lower edge of the “beak”). The upper edge is a curve of constant hohlraum temperature. The higher the temperature, the greater the risk of laser plasma interactions, so the probability of ignition decreases as you move up and to the left in the plot. The lower edge is a curve of constant “capsule margin”—our judgement of the likelihood a particular capsule will ignite assuming the hohlraum performs as expected. Both implosion velocity, which determines ignition temperature, and shell aspect ratio, which determines sensitivity to instabilities, enter into this judgement. The probability of ignition decreases downward in Figure 1-17 because the capsule margin decreases. The numeric scale (0.7, 1.0, 1.3) is highly nonlinear; the probability of ignition is quite low at the 0.7 curve and reasonably likely at the 1.3 curve.

Note that the 1.8-MJ, 500-TW pulse point lies in the middle of the laser performance band in Figure 1-17. This means that the baseline Haan pulse should be achievable if the optimistic values of the safe operating limits are obtained. The optimistic values assume that a Σ -B of four radians is tolerable. With our most pessimistic estimates of these operating limits (Σ -B = 3.6), the 1.8-MJ, 500-TW pulse is beyond the recommended operating limit. In that case, we would be forced to use a 1.65-MJ, 450-TW pulse, which has about 10% less ignition margin than the baseline.




Figure 1-16. Comparison of all NIF ICF temporally shaped (Haan) pulse operating limits. The limits are for controlling 3ω beam breakup (UV power less than 750 TW, shown in blue), 1ω beam breakup ($\Delta B \leq 1.8$, shown in red), injection energy less than 2 J (shown in yellow), focal spot size ($\Sigma B \leq 3.6 - 4$, shown in green). The limits are displayed in UV energy-power representation.




Figure 1-17. Overall NIF ICF temporally shaped (Haan) pulse operating limit. The limit is obtained by combining the five individual limits in Figure 1-16. The limit is shown in UV energy-power representation. The width of the operating band is due to uncertainties in the focal spot size limit.

1.3 Performance Risk Due to Low Amplifier Gain

In this section, we estimate the effects of amplifier gain variations on NIF performance. The range of gain values studied is based on the potential gain threats listed in Table 1-2. The various gain risks are expressed as reductions from the average amplifier gain coefficient, which has units of percent gain per centimeter of amplifier glass. The NIF baseline requirement is 5% per centimeter. This value is used in our NIF baseline system computer models. (Actually, the computer models are somewhat more complicated and account for variations in gain at different positions in the beam aperture and for different locations of the slab within the amplifier. However, when averaged over the entire beam aperture and over all 16 slabs in the system, the average gain coefficient is about 5% per centimeter, the same as the baseline requirement.)

Table 1-2. Potential NIF amplifier gain threats.

Gain risk	Reduction in gain coefficient (%/cm)
Blastshield AR coating degradation	-0.13
Flashlamp reflector loss	-0.05
Water in laser glass	-0.06
Thin laser glass from refinishing	-0.11
Power conditioning aging	-0.035
Flashlamp envelope transmission	-0.08
Even more flashlamp reflector loss	-0.05
Thin laser glass from vendor	-0.11
Reoptimize cavity for lower wavefront	-0.07

The effect of realizing all the gain risks listed in Table 1-2 is to reduce the gain coefficient by 0.7% per centimeter, an unlikely event. Gain loss can be recovered, however, by adding additional capacitors to the pulsed power system. The addition of four additional capacitors per module is estimated to increase the gain coefficient by 0.25% per centimeter at an additional cost of \$5M.

Calculations show that an upside of the uncertainty in gain is a possibility that the amplifier gain might actually be greater than the baseline value. Current computer estimates of the average gain coefficient predict a value of 5.15% per centimeter if none of the gain risks in Table 1-2 are realized.

So, for the purposes of this study, we considered five different gain values:

1. All gain risks occur. No additional capacitors. Average gain coefficient equals 4.46% per centimeter.
2. All gain risks occur. Add four capacitors. Average gain coefficient equals 4.71% per centimeter.
3. Baseline. No gain risks occur. No additional capacitors. Average gain coefficient is 5% per centimeter.

4. Current gain estimate. No gain risks occur. No additional capacitors. Average gain coefficient is 5.15% per centimeter.
5. Current gain estimate. No gain risks occur. Add four capacitors. Average gain coefficient of 5.4% per centimeter.

The results of the analysis are shown in Figures 1-18 to 1-20. Figures 1-18(a) and (b) show NIF system performance with temporally flat pulses, Figures 1-19(a) and (b) correspond to Gaussian pulses, and Figure 1-20 is for ICF indirect-drive Haan pulses. For flat and Gaussian pulses, the figures show the power and energy characteristics in two ways: energy vs pulse length and power vs energy. Pulse length is understood to be the full-width half maximum (FWHM) for the Gaussian pulse. In the case of the ICF Haan pulse, only the power vs energy graph is included, since pulse length is not unambiguously defined for the Haan pulse.

In each graph, the case with all gain risks realized is plotted in red, the baseline gain is in black, and the current best estimate for gain is plotted in blue. The cases with the additional four capacitors added are denoted with dashed lines. The performance curves cross in some regions. The overlapping regions are the parts of the curve where the reduced B-integral due to thin laser glass from refinishing is balanced by reduced gain.

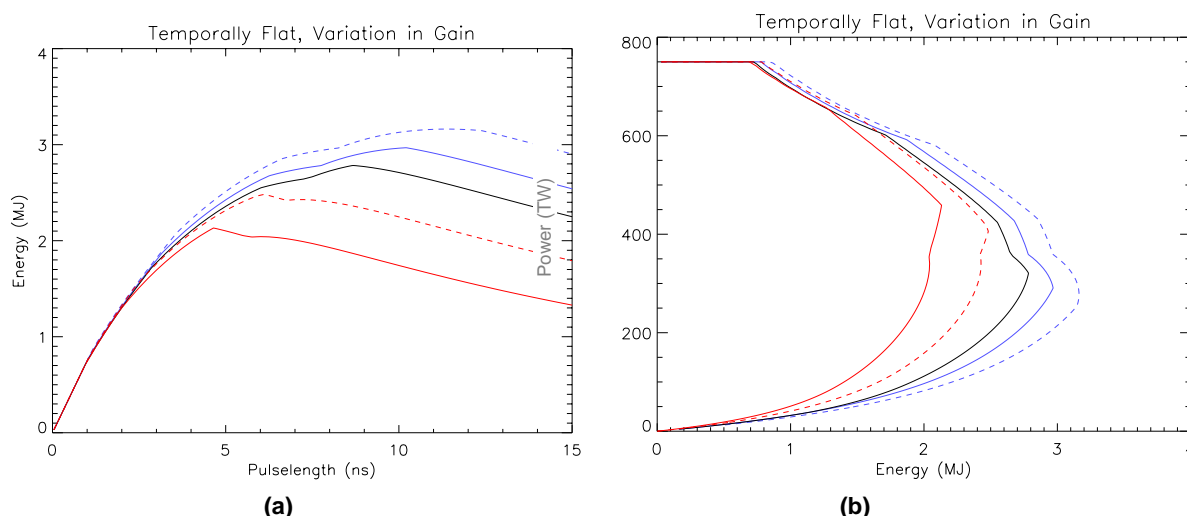


Figure 1-18. Sensitivity of NIF flat-pulse UV performance with respect to variations in amplifier gain. (a) System performance is shown in the UV energy pulse-length representation. Five different amplifier gains are shown. The NIF baseline amplifier gain (5% per cm) is shown in black. The expected gain if all amplifier risks are realized (4.46% per cm) is shown in solid red. The 1999 best estimate (5.15% per cm) is shown in solid blue. Performance increase with the addition of four capacitors to boost gain is shown by dashed lines. (b) The information is the same as in (a), but is redrawn to show UV power as a function of energy.

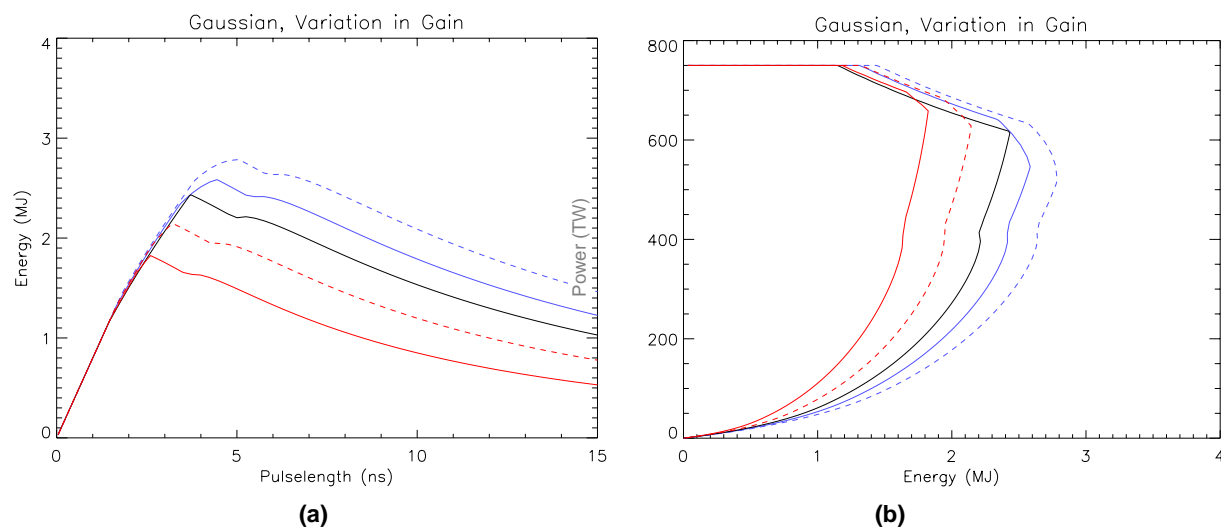


Figure 1-19. Sensitivity of NIF Gaussian-pulse UV performance with respect to variations in amplifier gain. (a) System performance is shown in the UV energy pulse-length representation. Five different amplifier gains are shown. The NIF baseline amplifier gain (5% per cm) is shown in black. The expected gain if all amplifier risks are realized (4.46% per cm) is shown in solid red. The 1999 best estimate (5.15% per cm) is shown in solid blue. Performance increase with the addition of four capacitors to boost gain is shown by dashed lines. (b) The information is the same as in (a), but is redrawn to show UV power as a function of energy.




Figure 1-20. Sensitivity of NIF ICF temporally shaped (Haan) pulse UV performance with respect to variations in amplifier gain. System performance is shown in the UV energy pulse-length representation. Five different amplifier gains are shown. The NIF baseline amplifier gain (5% per cm) is shown in black. The expected gain if all amplifier risks are realized (4.46% per cm) is shown in solid red. The 1999 best estimate (5.15% per cm) is shown in solid blue. Performance increase with the addition of four capacitors to boost gain is shown by dashed lines.

The simplest way to interpret the results is to compare the maximum system energy for the various gain cases. Table 1-3 summarizes the energy capabilities of each of the 15 cases considered (three pulse families with five gain cases each.) For each case the table lists the maximum 3ω system energy. Of course, reduced gain has little impact on the system's short-pulse high-power capabilities, and each of the cases is capable of achieving the 750-TW UV beam breakup limit. Therefore, gain risk is not much of an issue for meeting the high peak power, short pulse-length missions.

Table 1-3. Summary of possible energy capabilities for each of the proposed cases.

Gain Case	All Gain Risks	All Gain Risks + 4 Capacitors	Baseline	1999 Estimate	1999 Estimate + 4 Capacitors
Flat Pulse 3ω Energy (MJ)	2.13	2.48	2.78	2.97	3.16
Gaussian Pulse 3ω Energy (MJ)	1.82	2.14	2.43	2.58	2.78
Haan Pulse 3ω Energy (MJ)	1.38	1.73	2.06	2.27	2.54

1.4 Flowdown Analysis: 1ω Laser Performance and Front-end Requirements

As an extension of the previous results, we calculated the 1ω laser performance and the injection energy requirements for each of the five gain values discussed in Section 1.3. The 1ω energy entering the frequency conversion crystals and the injection energy were calculated for each of the three temporal pulse shapes (flat, Gaussian, and ICF indirect drive,) studied in the previous sections. For each pulse shape, the 1ω performance was calculated as limited by the three safe operating limits relevant to the main laser: $\Delta B \leq 1.8$, $\Sigma B \leq 4.0$, and injection energy of less than 2 J. Increasing the injection energy to a value exceeding 2 J will increase the output energy, but at a cost of pulse shape stability. Operating limits for 3ω beam breakup and UV damage are not included since they are intrinsically related to the 3ω section of the laser.

The $1\text{-}\mu\text{m}$ laser performance is summarized in Figures 1-21 to 1-23. Figures 1-21(a) and (b) show the system performance for temporally flat pulses expressed in terms of energy vs pulse length (Figure 1-21[a]) and peak power vs energy (1-21[b]). Figures 1-22(a) and (b) are the corresponding plots for Gaussian pulses. Finally, Figure 1-23 gives the $1\text{-}\mu\text{m}$ power vs energy performance using ICF shaped pulses.

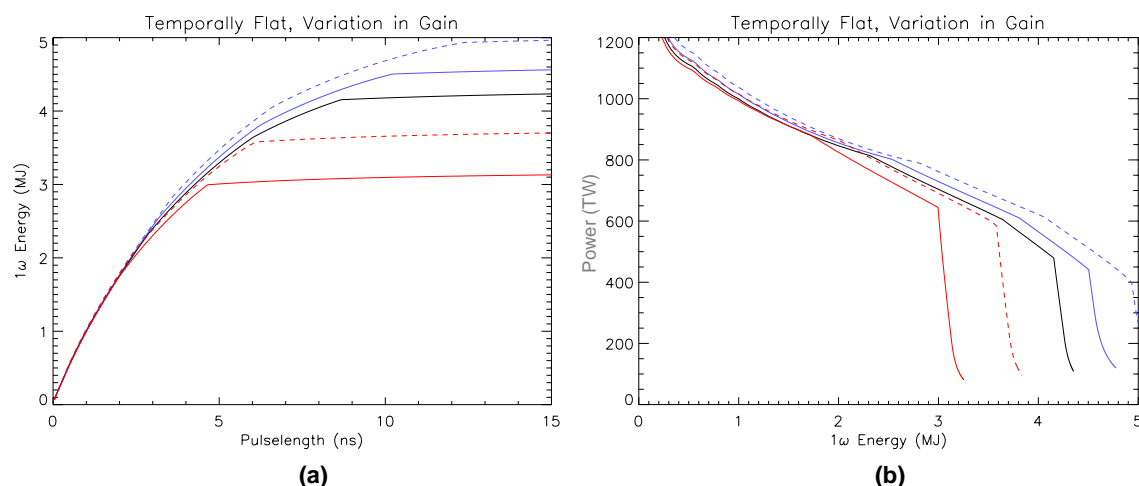


Figure 1-21. Sensitivity of NIF flat-pulse 1- μm performance with respect to variations in amplifier gain. (a) System performance is shown in the 1- μm energy pulse-length representation. Five different amplifier gains are shown. The NIF baseline amplifier gain (5% per cm) is shown in black. The expected gain if all amplifier risks are realized (4.46% per cm) is shown in solid red. The 1999 best estimate (5.15% per cm) is shown in solid blue. Performance increase with the addition of four capacitors to boost gain is shown by dashed lines. (b) The information is the same as in (a), but is redrawn to show 1- μm power as a function of energy.




Figure 1-22. Sensitivity of NIF Gaussian-pulse 1- μm performance with respect to variations in amplifier gain. (a) System performance is shown in the energy pulse-length representation. Five different amplifier gains are shown. The NIF baseline amplifier gain (5% per cm) is shown in black. The expected gain if all amplifier risks are realized (4.46% per cm) is shown in solid red. The 1999 best estimate (5.15% per cm) is shown in solid blue. Performance increase with the addition of four capacitors to boost gain is shown by dashed lines. (b) The information is the same as in (a), but is redrawn to show 1- μm power as a function of energy.




Figure 1-23. Sensitivity of NIF ICF temporally shaped (Haan) pulse 1- μ m performance with respect to variations in amplifier gain. System performance is shown in the 1- μ m energy pulse-length representation. Five different amplifier gains are shown. The NIF baseline amplifier gain (5% per cm) is shown in black. The expected gain if all amplifier risks are realized (4.46% per cm) is shown in solid red. The 1999 best estimate (5.15% per cm) is shown in solid blue. Performance increase with the addition of four capacitors to boost gain is shown by dashed lines.

Again, a simple figure of merit for judging the system performance for the various gain cases is the maximum 1- μ m energy extractable with 2 J. As can be seen by comparing Figures 1-21 to 1-23, the maximum extractable 1 ω energy does not depend strongly on the type of pulse propagated. As a summary, Table 1-4 shows the maximum 1- μ m energy for each of the five gain cases studied using temporally flat pulses. The total spread in extractable energy is almost 2 MJ between the worst case with all gain risks realized and the best case with 1999 gain estimates, no gain risks realized and four additional capacitors.

Table 1-4. Summary of maximum 1- μ m energy of each of the five gain cases.

Gain Case	All Gain Risks	All Gain Risks + 4 Capacitors	Baseline	1999 Estimate	1999 Estimate + 4 Capacitors
Max 1 ω Energy (MJ)	3.25	3.85	4.45	4.80	5.20

Finally, Figures 1-24 to 1-26 show the injection energy requirements for each of the five gain values studied. Injection energy per beamline is plotted as a function of total 192-beam 1 ω energy for each pulse: flat, Gaussian, and ICF indirect drive. Again,

injection energy requirements do not depend strongly on pulse shape. Therefore, the results in Figures 1-24 to 1-26 are similar.




Figure 1-24. Sensitivity of NIF flat-pulse injection energy requirements with respect to variations in amplifier gain. Injection energy is plotted as a function of 1- μ m output energy at the frequency converter crystals for five different amplifier gain values. The NIF baseline amplifier gain (5% per cm) is shown in black. The expected gain if all amplifier risks are realized (4.46% per cm) is shown in solid red. The 1999 best estimate (5.15% per cm) is shown in solid blue. Performance increase with the addition of four capacitors to boost gain is shown by dashed lines.




Figure 1-25. Sensitivity of NIF Gaussian-pulse injection energy requirements with respect to variations in amplifier gain. Injection energy is plotted as a function of 1- μm output energy at the frequency converter crystals for five different amplifier gain values. The NIF baseline amplifier gain (5% per cm) is shown in black. The expected gain if all amplifier risks are realized (4.46% per cm) is shown in solid red. The 1999 best estimate (5.15% per cm) is shown in solid blue. Performance increase with the addition of four capacitors to boost gain is shown by dashed lines.




Figure 1-26. Sensitivity of NIF ICF temporally shaped (Haan) pulse injection energy requirements with respect to variations in amplifier gain. Injection energy is plotted as a function of 1- μm output energy at the frequency converter crystals for five different amplifier gain values. The NIF baseline amplifier gain (5% per cm) is shown in black. The expected gain if all amplifier risks are realized (4.46% per cm) is shown in solid red. The 1999 best estimate (5.15% per cm) is shown in solid blue. Performance increase with the addition of four capacitors to boost gain is shown by dashed lines.

References

1. S. W. Haney, W. H. Williams, et al., “NIF Energetics under reduced slab count/reduced PAM size operation,” NIF-0000582, 1996.
2. O. Jones, et al., “The NIF’s power and energy ratings for ICF-shaped temporal pulses,” *Supplement to Proc. SPIE* **3942**, 49–54 (1998), C. Widmayer, et al., “The NIF’s power and energy ratings for flat in time pulses,” *Supplement to Proc. SPIE* **3942**, 11–21 (1998).
3. W. Williams et al., “NIF’s basic focal spot for temporally flat pulses,” *Supplement to Proc. SPIE*, **3942**, 22–38, (1998).
4. The Filamentation Working Group, “Beam filamentation on Beamlet and Nova,” NIF-LLNL-96-476, (1996).
5. W. Williams, et al., “Modeling of self-focusing experiments by beam propagation codes,” *ICF Annual Report*, UCRL-LR-105821-96, 7–21, (1996).

2.0 INJECTION LASER SYSTEM (ILS).....	2-1
2.1 Introduction.....	2-1
2.2 MOR Overview	2-3
2.3 MOR Procurement Strategy	2-8
2.4 PAM Overview.....	2-9
2.5 MOR/PAM Prototype Integrated System Testing	2-11
2.6 PAM operation for ICF and non-ICF Missions	2-15
2.7 PAM 50 Shot Pointing and Far-Field Stability Experiment.....	2-17
2.8 PAM Procurement Strategy	2-21
2.9 PABTS System Overview	2-22
2.10 Procurement Strategy	2-26
Building 381 Hibay First-Article PABTS/PASS Integrated Testing.....	2-27

2.0 Injection Laser System (ILS)

Mark Hermann, Isaac Bass, Ernie Dragon, Don Jedlovec, Russ Wilcox,
Mark Henesian and Dick Hackel

2.1 Introduction

The purpose of this report is to update the status and performance of the injection laser system (ILS). This report summarizes the current technology, performance, and the procurement status of the master oscillator room (MOR) system, the preamplifier module (PAM) system, and the preamplifier beam transport system (PABTS). In addition, planned ILS front-end integrated system test (FEIST) experiments in Building 381 and PABTS first-article integration experiments in the Building 381 high bay will be described.

The ILS consists of three subsystems:

1. The MOR produces 48 temporally shaped pulses with ~ 0.5 to 2-nJ pulse energy employing a fiber-based laser system.
2. The PAM amplifies the MOR-shaped pulse to as much as 16 J (10- to 20-ns pulse) with square-pulse distortion (SPD) of less than 2.3 and a spatially shaped beam to compensate for spatial gain nonuniformity in the NIF amplifiers. There are 48 PAMs in the NIF.
3. The PABTS splits the PAM output into four beams with a maximum energy of 3 J (10- to 20-ns pulse) per beamline. The PABTS also provides isolation to the PAM and MOR from light that might be back-reflected into them from the main amplifiers. There are 48 PABTS quads in the NIF.




Table 2-1 summarizes the major performance criteria of the ILS system at the injection window of the transport spatial filter, which is the interface with the main laser. Tables 2-2, 2-4, and 2-8 are requirements of the MOR, PAM, and PABTS obtained by flowing down Table 2-1 criteria to the ILS subsystems.

Table 2-1. ILS Specifications at Injection Window.*

Parameter	Proposed Requirement	SSDR**
Wavelength	1053.00 nm	same
Wavelength stability	± 10 picometers (pm)	need to add
Energy and adjustment precision	100 mJ to 3 J , 0.1% precision	same
Pulse repetition frequency	1 per 10 minutes	modify
Pulse length	0.5–20 ns, “arbitrary” shape	same
Pulse dynamic range	$\geq 125:1$ @ 3 J	same
Pulse timing jitter	≤ 15 ps rms	same
Timing adjustment: Quad—remote adjustment	≤ 30 ns minus pulse length***	same
Quad—gross adjustment (not remote)	≤ 100 ns quad-to-quad	same
Beamline—remote adjustment	± 125 ps (< 1 ps adjustability)	same
Beamline—gross adjustment (not remote)	± 5 ns (< 30 ps adjustability)	same
Square pulse distortion	< 2.3 @ 3 J	same
ASE @ 1.5 J per beamline	< 5 kW in 3 ns	modify
Prepulse @ 1.5 J	< 5 kW in 20 ns, < 20 μ J	modify
Spatial beam shape contrast (maximum) & deviation	2:1 & < 10%	modify
Near-field contrast (@ RP10 after spatial filtering)	< 7.5% rms	need to add
Wavefront ($\leq 3^{\text{rd}}$ -order contribution)	< 1.5 waves	same
Divergence (radius of encircled energy) [After 1.5 waves of < 3 rd -order correction] [Ref. 372 \times 372 mm beam]	> 80% ≤ 6 microradians (μ rad) < 0.5% ≥ 33 μ rad	equivalent equivalent
Pointing stability [Ref. 372 \times 372 mm beam]	< 1.2 μ rad rms	modify
Energy stability, power balance	< 3% rms, < 3% rms in any 2-ns interval	modify
SSD bandwidth	3Å @ 17 GHz	same
SBS bandwidth	1Å @ 3 GHz	same
Beam size	45 \times 45 mm ² (+0, -0.2 mm)	same
Back reflection isolation	> 30 dB & ≤ 5 J/cm ² fluence	modify
Beam rotation (maximum)	< 5 mrad	need to add
FM-AM conversion	< 15%	need to add

*All requirements are for 1Å bandwidth undispersed and the Haan baseline indirect-drive ICF pulse shape only.

** We are in the process of modifying, adding, or clarifying our SSDRs. The status of the listed requirements with respect to the current SSDR is as follows:

- Same = as stated in the Software Subsystem Design Requirements (SSDR).
- Modify = need to change existing SSDR to meet NIF requirements.
- Equivalent = need to restate but equivalent to SSDR.
- Need to add = no SSDR but requirement is needed for the NIF.

***The maximum quad timing remote adjustment is the difference between the input pulse length (30 ns) to the MOR amplitude modulator chassis and the derived pulse length.

2.2 MOR Overview

The MOR (see Figure 2-1) is a fiber laser distribution system based on telecommunications technology employing an LLNL-designed and built chassis and commercially built components and chassis based on LLNL requirements. The MOR subsystem performance requirements, which are obtained by flowing down the ILS interface requirements to the MOR are summarized in Table 2-2. These requirements are at the beam-forming aperture location in the PAM. A summary of the MOR component specifications is given in Table 2-3.

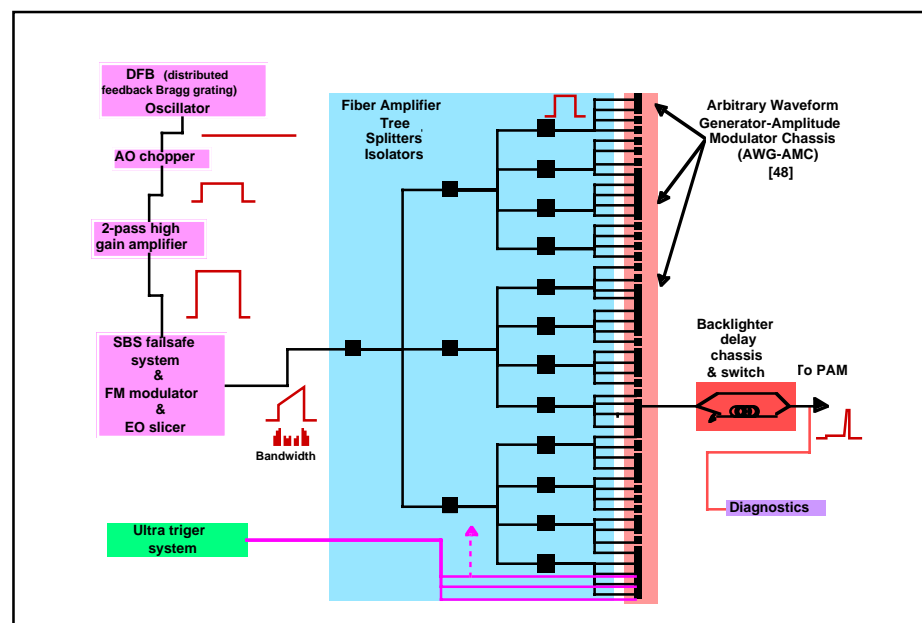


Figure 2-1. The MOR system produces 48 shaped pulses with ≤ 12 ps timing jitter.

Table 2-2. MOR subsystem requirements at input to PAM.

Parameter	NIF Spec	Demonstrated
Wavelength	1053.00 nm	1053.00 nm
Wavelength stability	± 10 pm	$< \pm 1.5$ pm
Output power at PAM	100 mW	100 mW
Energy stability	5% rms	5.5% rms
Power balance	2% rms in 2 ns	Not measured
Repetition rate	960 Hz	960 Hz
Prepulse extinction	$\geq 10^5$	Not measured
Pulse shape	Variable	Square
Pulse length	0.2–20 ns	4 ns
Peak/foot contrast	275:1	Not measured
Pulse timing jitter	≤ 12 ps	Not measured
Phase modulation (SSD)	1 Å @ 3 GHz 3 Å @ 17 GHz	3 Å @ 3 GHz
FM-AM conversion	10% rms	Not measured

Table 2-3. MOR component specifications and status of tests.

Component	Parameter	NIF specification	Measured
Distributed feedback (DFB) oscillator	Power	≥ 20 mW	21 mW
	Power stability	$\leq 1\%$ rms	$< 1.4\%$ rms
	Wavelength stability	$\leq \pm 10$ pm	$< \pm 1.5$ pm
2-pass high gain amp & AO chopper	Pulse length	100 ns	100 ns
	Gain	200	350
	Power out	1–2 W	1 W
	PRF	960 Hz	960 Hz
	Power stability	$\leq 1\%$ rms	$\leq 1.4\%$
	ASE content	–50 dB ASE @ 3 nJ	–58 dB ASE @ 3 nJ
SBS failsafe, Phase modulator, EO pulse slicer	Pulse length	30 ns	Not measured
	PM (SBS)	1 Å @ 3 GHz	Not measured
	PM (SSD)	3 Å @ 17 GHz	Not measured
	Pulse shape	Shaped, 2.5:1 contrast	Shaped, $> 2.5:1$
	SBS failsafe	RAM statistics	Not measured
Fiber amplifier tree	Small-signal gain	32 per amplifier	> 80 per amplifier
	Power into arbitrary waveform generator (AWG)-amplitude modulator chassis (AMC)	3 W	Not measured
AMC	Short pulse mode	≤ 200 ps	Not measured
	Shaped pulse mode	0.5–20 ns	Not measured
		250 ps resolution	Not measured
	Dynamic range	275: 1	Not measured
	Rise time	≤ 200 ps	
PZ fiber	Polarization extinction & 1053 nm absorption		
Type A		≥ 5 dB/m, ≤ 0.2 dB/m	> 7 dB/m, ~ 0.023 dB/m
Type B		≥ 2 dB/m, ≤ 0.02 dB/m	4.4 dB/m, 0.0053 dB/m
MOR integrated system	Power (@ PAM)	100 mW	Not measured
	FM-AM conversion	7.5%	Not measured
	Power balance	$\leq 2\%$ rms in 2 ns	Not measured
	Energy stability	$\leq 5\%$ rms	5.5%

A single, continuous wave (CW) master oscillator employing a fiber laser with DFB Bragg grating reflectors is the seed laser for the entire NIF. Consequently, it needs to be robust and reliably produce low-noise CW power in a single longitudinal mode to prevent optical damage in the entire NIF chain. The DFB laser by design does not support more than one longitudinal mode. The DFB oscillator gain is provided by 980-nm, diode-pumped, Ytterbium-doped silica fiber similar to the Erbium-doped gain fiber commonly employed in the telecommunications industry. The NIF oscillator operates at 1053 nm. A commercial vendor supplies the DFB fiber oscillator, which LLNL integrates into a chassis with wavelength stabilization and isolation. Figure 2-2 shows the measured relaxation oscillations of the first-article NIF DFB laser oscillator

producing less than 0.4% power fluctuations at ~20 mW, which exceeds our specifications. Figure 2-3 displays the wavelength stability (± 1 pm) of the oscillator, which greatly exceeds the ± 10 -pm specification. We will temperature tune the DFB oscillator to operate at 1053.0 nm. The oscillator itself is single-mode fiber as opposed to the polarizing (PZ) fiber employed in the majority of the MOR system.




Figure 2-2. Short-term power stability of <0.4% peak-to-peak with no relaxation oscillations has been measured for first-article DFB fiber oscillator. (Note that vertical axis scale is ~0.16% of mean.)

The CW output is chopped by an acousto-optic modulator into 100-ns pulses at 960 Hz. These pulses are amplified by an LLNL double-pass fiber amplifier. The double-pass amplifier has a gain in excess of 300 and removes amplified spontaneous emission (ASE) noise by spectral filtering. The measured input and output energies of the double-pass amplifier are displayed in Figure 2-4. The measured stability is 1.4% rms. Much of the long-term fluctuation is due to temperature-induced polarization changes at the oscillator's single-mode to PZ-fiber connection. We plan on improving the stability of the oscillator/double-pass amplifier by securing the single-mode fiber, better temperature control in the racks, and possibly an electronic feedback control system. The oscillator and double-pass amplifier first-article chassis have been successfully built, tested, and integrated into equipment racks in the MOR FEIST lab in Building 381. A photograph of the oscillator and double-pass amplifier chassis in the equipment rack is shown in Figure 2-5.




Figure 2-3. Long-term DFB first-article fiber oscillator output wavelength stability. The periodic structure is due to temperature cycling in the lab. The measured wavelength stability of ± 1 pm exceeds the required stability of ± 10 pm.




Figure 2-4. Input energy and output energy of the first-article double-pass amplifier (100-ns full-width at half maximum [FWHM] pulse). We are currently adding active feedback control to reduce the 1.4% rms energy stability to 1%.



Figure 2-5. Photograph of first-article DFB oscillator, double-pass amplifier, and PZ-fiber high-gain amplifier (left-to-right).

The Stimulated Brillouin Scattering (SBS) failsafe subsystem applies 1-Å bandwidth at 3-GHz modulation to prevent SBS-induced damage in 40-cm optics. It consists of a diagnostics package, fiber delay line, and electro-optic (EO) modulator identical to that employed in the amplitude modulator chassis. The diagnostics measure the bandwidth. If proper bandwidth is not detected on a given pulse, then that pulse is not permitted to propagate through the NIF system by the EO modulator. A fiber delay line between the diagnostics and EO modulator in the SBS failsafe system permits the diagnostics and electronics time to trigger the EO modulator before the pulse arrives. In addition, the EO modulator reduces the pulse length to 30 ns and shapes it to compensate for square-pulse distortion in the fiber amplifiers. The failsafe prototype modulator and electronics have been successfully demonstrated. In addition, 3 to 5 Å of bandwidth at 17-GHz modulation can be applied for focal spot smoothing by spectral dispersion (SSD). The EO modulator chassis and its pulse-shaping drivers have been built, tested, and integrated into equipment racks in the MOR FEIST lab in Building 381. The SBS failsafe system will be tested this year continuously for two months at 1 kHz to demonstrate the required reliability for the NIF system.

A fiber amplifier tree splits and amplifies the 30-ns pulse from the SBS-failsafe subsystem into 48 pulses. The LLNL-designed production amplifier chassis including the splitters will be built by a commercial vendor. First-article chassis employing PZ fiber are being built and tested at LLNL. A photograph of the first-article amplifier chassis is shown in Figure 2-5. Earlier prototypes using polarization maintaining (PM) fiber suffered from thermal-drift-induced energy fluctuations and FM-AM conversion due to polarization dispersion in PM fibers.

The amplified pulses are then shaped in the amplitude modulator chassis (AMC). The AMC includes an arbitrary waveform generator for generating Haan-type ICF pulse shapes, or alternatively, 200-ps timing pulses. The chassis is being developed and built by a commercial company. The first chassis is scheduled for delivery in February 2000. One primary issue with the MOR is that we have not generated a shaped pulse (Haan) with power balance in the foot that meets requirements. However, a 4-ns prototype

board (1 of 8 used in the AMC) was successfully tested. Power balance and pulse-shaping measurements will be part of our integration activities in Building 381.

The AMCs are triggered optically by the ultratrigger system to meet the 12-ps rms MOR timing allocation of the 30-ps overall NIF requirement. The ultratrigger will be built and tested next year.

The temporally modulated pulse is then delivered to the PAM via a set of forty-eight 80- to 100-m-long fibers. Most of the fiber in the MOR and fiber distribution system is PZ fiber. PZ fiber as opposed to PM fiber is used to minimize large fluctuations in energy due to polarization mixing caused by small thermal drifts or through stressed fiber connectors. PZ fiber is also required to minimize FM-AM conversion that could cause laser spiking-induced damage, power balance errors, and a reduced conversion efficiency to 3ω .

Another technical issue for the MOR system is PZ fiber. Because PZ fiber is in a high-stress state, it is difficult to splice and connectorize. However, we have made considerable progress in splicing the fiber, and our component vendors are now able to incorporate PZ fiber into MOR components. We use two types of PZ fiber in our baseline design, a type-A and a type-B. The type-A has a high polarization extinction and moderate loss. It is used for the many short fiber jumpers between chassis in the MOR to extinguish the depolarization effects caused by the connectors. The type-B has a lower polarization extinction and low loss. It is used for the long fibers to the PAMs to minimize losses.

2.3 MOR Procurement Strategy

The MOR system in the NIF will be composed of roughly 230 electronic and electro-optic line replaceable units (LRUs) built in electronic-chassis configurations. They will be mounted in electronic racks. There are approximately 40 different types of these chassis. Roughly 15 are directly available as commercial, catalogue products. All the remaining chassis, with the exception of three types (see below), will be designed and built by LLNL personnel.

The MOR chassis to be built at LLNL are typically in quantities less than five of a given type, or they are relatively simple in design. These chassis will be built in FY00 and the early part of FY01.

The three chassis to be built by outside vendors are the AMC, the dual-gate (backlighter) chassis, and the fiber amplifier chassis. There are a large number of them, and they represent a significant fraction of the MOR procurement budget. A total of 48 AMCs will be designed and built by a vendor of high-speed electronics. Half of them will be procured in FY02 and FY03 for the first 96 NIF beams, and the remainder in FY05 and FY06. The procurement of the 25 dual-switch chassis is similarly scheduled.

An outside vendor expert in telecommunications fiber optics technology will build the 18 fiber amplifier chassis on a build-to-print basis. The schedule calls for four in FY00, seven in FY02, and seven in FY05.

In addition to the various chassis, roughly 20,000 meters of special, PZ fiber must be procured for the NIF. The first 5000 meters was ordered in FY99. The remainder must be ordered in FY02 to pull all the fiber cables from the MOR to the laser bays at the same time.

2.4 PAM Overview

The PAM is designed to amplify the 0.5–2-nJ shaped pulse from the MOR up to 16 J (10- to 20-ns pulse) with a temporal square pulse distortion (SPD) of less than 2.3. In addition, the beam must be spatially shaped to compensate for gain nonuniformity in the main amplifiers. To date, all of our experimental efforts have focused on verifying key NIF performance requirements using an engineering prototype of the PAM. The engineering prototype is based on the same LRU package envelope as the current baseline production PAMs. The NIF requirements for the PAM and recent performance measurements of the engineering prototype are shown in Table 2-4. A conceptual layout of the PAM is presented in Figure 2.6, and corresponding photographs are shown in Figure 2.7.

Table 2-4. PAM specifications and prototype performance at output.

Performance Parameter	NIF Specification	Demonstrated
Energy & stability	30 mJ to 16 J @ 3% rms	16 J @ 2.6% rms
PRF (@ 16 J)	1 per 20 minutes	1 pulse per 10 min
Regen PRF	1 Hz	1 Hz
Beam size	30 × 30 mm ²	30 × 30 mm ²
Spatial beam shape	2:1	2:1
Pulse length	0.2 – 20 ns	8 ns
Pulse dynamic range	125:1 @ 16 J	Not measured
SPD square pulse distortion	≤ 2.3 @ 16 J	2.3 @ 17 J
ASE @ 16 J @ PAM output	≤ 96 kW in <5 ns	96 kW
Prepulse @ 16 J	≤ 600 W in 5–200 ns	80 kW
Near field contrast uniformity	≤ 6%	7%
Wavefront (≤ 3 rd order)	≤ 1.2 waves	1 wave
Divergence (encircled energy radius) referenced to 372 mm beam. [After <3 rd order correction]	> 80% in 4.8 μrad > 95% in 10.7 μrad < 0.5% outside 208 μrad	> 80% in 4.6 μrad > 95% in 12.6 μrad In progress
Pointing stability (372 mm beam)	≤ 0.96 μrad rms	2.6 μrad rms
Beam rotation	5 mrad rms	Not measured
Power balance (PAM+MOR)	< 3% rms	Not measured
Shot-to-shot timing jitter (PAM+MOR)	8 ps rms	Not measured
FM-AM conversion (PAM+MOR)	10% rms	Not measured

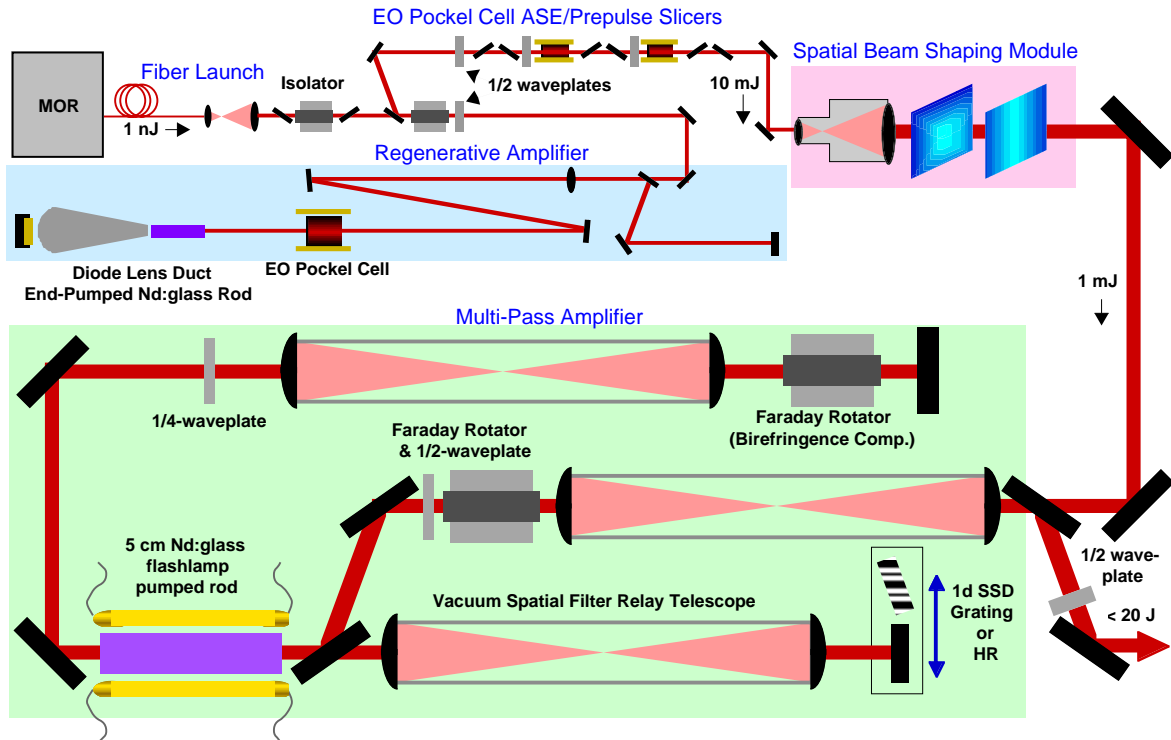


Figure 2-6. The PAM amplifies the ~ 1 nJ MOR input to ~ 20 J and spatially shapes the beam to compensate for gain nonuniformities in the main amplifiers.




Figure 2-7. PAM Engineering prototype showing fiber launch, regenerative amplifier, beam-shaping module, 5-cm flashlamp-pumped amplifier, and multipass amplifier hardware.

The MOR pulse is delivered to the 48 PAMs via a PZ fiber distribution system. The pulse is first amplified in a regenerative amplifier (regen) to ~15 mJ using a diode-end-pumped (DPSSL) Nd:glass rod amplifier. The diodes are actively temperature-stabilized to eliminate energy drift, which was observed in early engineering prototype experiments. An important change to the original baseline design was changing the regen operating mode from unsaturated to saturated gain operation. Since the optical pulse length of the regen is longer than the pulse, the entire pulse is amplified before the next round trip. Since the gain is low, many round trips are required, and with each subsequent round trip, the gain is gradually reduced. This results in minimal temporal pulse distortion. Saturated operation reduces the regen cost because: (1) 20% fewer diodes are required and (2) the costly feed-forward control loop required to meet the stringent energy stability requirement is eliminated. The advantage of this system modification has been demonstrated with the PAM prototype. The gain of the DPSSL rod operates at a fixed point limited by the number of diodes. The number of round trips is also fixed, limited by the pulse length of the Pockels cell's high-voltage pulser.

Two Pockels cells and polarizers located after the regen are used to reduce the regen energy leakage and ASE prepulses. The ASE and prepulse powers reported in Table 2-4, which do not meet NIF requirements, were measured with a single Pockels cell and will be reduced by an additional 30 dB by the addition of a second Pockels cell. A waveplate-polarizer combination is used to adjust the input energy into the 4-pass amplifier.

The regen pulse is shaped in the spatial beam-shaping module (BSM) before injection into the MPA. The BSM contains spatial filters and a serrated aperture to convert the TEM₀₀ Gaussian beam to a square-beam with the spatial profile needed to compensate for gain nonuniformity in the 40-cm aperture amplifiers. The BSM transmission is ~10% producing ~1 to 2 mJ (>3-ns pulse) input into the 4-pass amplifier.

The 4-pass amplifier increases the beam energy to 16 J (10- to 20-ns pulses) with a SPD of 2.3. The 4-pass amplifier contains a 5-cm rod pumped by six flashlamps. The gain of the amplifier is fixed at the optimized flashlamp charge voltage, which minimizes gain fluctuations. A Faraday rotator corrects the birefringence produced by the rod and increases parasitic hold-off due to birefringence. Experiments indicate that at a rod gain of >18 (>20-J operating point) and a repetition rate faster than one shot per 40 min, a cavity Faraday rotator is needed to correct rod birefringence and hold off parasitic lasing in the 4-pass system.

A controllable waveplate after the 4-pass is used to adjust the injected energy into PABTS. Although there are controllable waveplates in the PABTS to adjust individual beamline fractions, the total quad energy can only be adjusted in the PAM.

2.5 MOR/PAM Prototype Integrated System Testing

We completed a series of integrated system tests measuring wavefront, energy and pointing stability, beam shaping, energetics, and square-pulse distortion. Figure 2-8

shows the energy stability measurements at the regen output versus the MOR energy input. The short-term MOR energy drift is 1.2%; however, the long-term energy drift of 5.5% is due to thermal cycling in the lab of the PM fiber in the fiber amplifiers (PZ fiber first-article components should reduce energy drift). As can be seen, the saturated regen reduces the 5.5% MOR energy fluctuations to 1.6%. We have measured PAM's energy stability at 17 J (8-ns pulse) output of 2.6% rms, which meets the NIF specification flow down to the PAM output. Figure 2-9 displays the pulse shape of the MOR input, the regen output, and the PAM. At the required 17 J, the SPD at output of the PAM is 2.3 (with a spatially shaped beam). The maximum output energy (Figure 2-9) that we have extracted is ~34 J, but at a higher SPD (~3). Figure 2-10 shows the spatially shaped output beam at 17 J.




Figure 2-8. The saturated gain mode of operation in the regen reduces the long-term (2 hr) NIF energy drift from 5.5% to 1.6% (left). The short time (5 min) stability of the MOR is 1.2%, and the regen remains at 1.6%. The requirement for the PAM is 3%.




Figure 2-9. Square-pulse distortion at 17 J (shaped) and demonstrated energy versus SPD for MOR output, regen output (1.33) and PAM output (2.3) meets the requirement of 2.3 at 16 J.

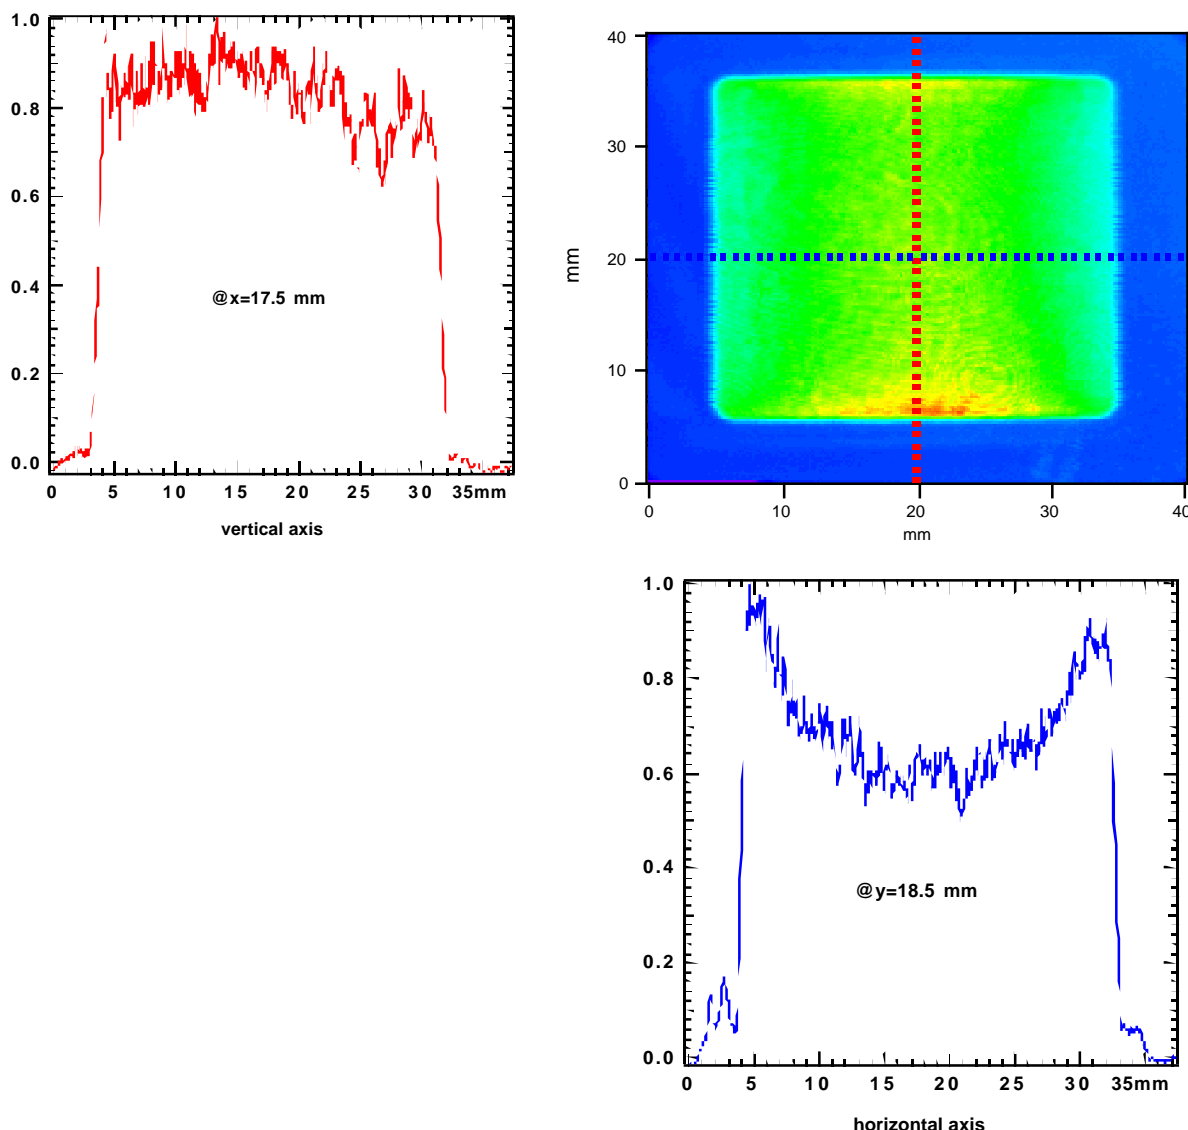



Figure 2-10. Shaped beam profile at PAM output.

Figure 2-11 is the measured and numerically smoothed wavefront of the beam. The peak-to-valley wavefront error is ~ 1 wave. Most of the low spatial wavelength error will be removed when the relay lenses are redesigned. Currently, all of the aspherical corrective power of the six 4-pass lenses is performed in a single lens to reduce cost. Since the beam propagates off-axis and with a tilt with respect to the lenses to eliminate on-axis parasitics and pencil beams, the beam's wavefront is distorted by the highly aspheric lens. The new design will distribute the aspherical correction among several of the 4-pass lenses and modify the propagation path and lens tilt. The measured focal spot and the focal spot with third-order aberrations removed are shown in Figure 2-12. Up to 0.9 waves of third-order aberration is budgeted to the PAM for removal by the deformable mirror. Figure 2-13 compares the encircled energy of the corrected focal spot to that of a perfect focal spot.

17 J measured wavefront with tilt removed



Zernike polynomial fit to measured wavefront




Figure 2-11. Measured and numerically smoothed wavefront of 27-J PAM beam.

Measured far field



Predicted far field
(3rd order corrected with
deformable mirror)

Figure 2-12. Measured focal spot (right) at 27 J and corrected estimate based on removing ~ 1 wave of third-order aberrations. The angles shown on the horizontal and vertical axes have been scaled to correspond to the beam size of 372 mm. The circles correspond to 50%, 80%, and 95% encircled energy. Eighty percent of the energy is enclosed by a circle with a radius of 4.6 microradians.




Figure 2-13. Encircled energy fraction of corrected focal spot (3rd-order corrections removed) from Figure 2-12 (solid) and perfect wavefront (dashed), scaled to a beam size of 372 mm.

2.6 PAM operation for ICF and non-ICF Missions

We have completed the energetic analysis of the integrated MOR and PAM for various NIF missions with pulse lengths ranging from 200 ps to 20 ns. Since the ILS system design was optimized for 1- to 3-J operation, there is concern that the energy stability for other non-ICF missions requiring short-pulse lengths with less energy may reduce system performance.

The number of round trips and gain of the DPSSL regen amplifier is optimized to saturate (saturated mode of operation) the gain for ICF (>3 ns, 1 J) pulses. Saturated regen operation greatly improves energy stability of the PAM. The drawback is that for short pulses (<1 ns), the regen does not completely saturate, resulting in an increase in energy fluctuations out of the regen. In addition, the 4-pass amplifier produces an increase in energy fluctuations since the gain saturates less for shorter, less energetic pulses. Injecting too much energy into the 4-pass in a short 200-ps pulse could damage the optics.

Table 2-5. ILS energetics for various pulse lengths.

Pulse length (ns)	10	4*	1	0.2
MOR energy into regen (nJ)	2.50	1.00	0.25	0.05
Regen output energy (mJ)	8.84	9.43	6.35	2.01
Power limited 4-pass input energy (mJ)**	0.88	0.80	0.20	0.04
4-pass output energy (J)	9.32	8.60	2.52	0.53
Maximum beamline injected energy (J)	1.63	1.51	0.44	0.09
Minimum beamline injected energy (J)	0.08	0.08	0.02	0.005

*Haan equivalent pulse

**After beam shaping (10% transmission) and energy control

We have completed an analysis of energy stability for various pulse lengths. The peak-peak (p-p) energy fluctuations were calculated with a Frantz-Nodvik model. The input instabilities for the model are based on component specifications and agree with experimental measurements and hardware specifications. Specifically, the measured regen diode pump stability is 1% p-p (0.25% rms), the 4-pass amplifier flashlamp pump stability is 1% p-p (0.25% rms), and the MOR energy stability is 20% p-p (5% rms). The calculated energy fluctuations of the regen and PAM output are presented in Table 2-6. The measured energy stability for an 8 ns pulse is 5.7% p-p (1.3% rms) for the regen output and 12% p-p (2.7% rms) for the PAM output, which agrees with the modeling predictions for 4-ns and 10-ns pulses. As can be seen in Table 2-6, the p-p energy fluctuation for 200 ps is predicted to be >60% (16% rms). Figure 2-14 shows measured shot-to-shot energy stability for MOR, regen, and PAM output.

Table 2-6. Calculated energy fluctuations for various pulse lengths.

Pulse length (ns)	10	4*	1	0.2
Regen output stability (peak-peak)	6%	6%	29%	47%
Regen output stability (rms assume 4 σ)	1.5%	1.5%	7%	22%
ILS output energy stability (peak-peak)	13%	13%	38%	64%
ILS output energy stability (rms assume 4 σ)	3%	3%	9.5%	16%

* Haan equivalent pulse




Figure 2-14. Shot-to-shot energy output of MOR, regen, and PAM. Sigma is rms, and p-p is peak-to-peak fluctuation.

We plan to perform energy stability measurements at various energies and pulse lengths corresponding to different NIF missions as part of FY00 front-end integrated system tests. These experiments may demonstrate the need for design upgrades of the ILS. We have identified a possible modification of the MOR system to reduce energy fluctuations for short pulses. Specifically, the amplitude modulator chassis will be modified to produce a pulse shape consisting of two parts: (1) the lead pulse will be the short (<3 ns) pulse injected into the laser system and (2) a 5-ns post pulse delayed 3–4 ns. The post pulse saturates the regen, reducing fluctuations. The post pulse is removed by switching the Pockels cells used for prepulse suppression. We plan on evaluating design modifications and implementing engineering changes if required.

2.7 PAM 50 Shot Pointing and Far-Field Stability Experiment

We completed a 50-shot 17 J experimental campaign that lasted approximately 8 hours duration. The PAM was fired once every 10 minutes. The purpose of the experiments was to measure both shot-to-shot and long-term energy, pointing, and beam quality stability. The measured energy stability of 2.6% meets the required 3% specification. The experiments did reveal several problems with the PAM leading to higher than expected pointing and far-field beam instabilities. During the experiments, we did not have capability to measure wavefront. Instead, we relied on the far-field beam profile as a metric for stability. If the far-field profile and encircled energy distribution for low-divergence angles fluctuates from shot-to-shot, it may be difficult for the deformable mirror to correct low-order aberrations because a portion of the low-order aberrations may be random. The Strehl ratio is defined as the peak irradiance at centroid of the measured far field divided by the peak irradiance at far-field centroid of a beam having the same near-field distribution, but a flat phase front. The Strehl ratio is a particularly sensitive metric of low-order aberrations. Table 2-7 summarizes the pointing and far-field fluctuations.

Table 2-7. PAM 50-shot run far-field and pointing fluctuations.

Measured Quantity ¹	Average	Fluctuation (rms)	Peak-to-peak
Divergence @ 50% ²	4.0 μrad	0.72 μrad	3.37
Divergence @ 80% ²	7.6 μrad	0.92 μrad	4.41
Divergence @ 90% ²	12.7 μrad	1.38 μrad	8.60
Divergence @ 95% ²	27.3 μrad	4.26 μrad	26.49
Times diffraction limit	2.91	0.33	1.72
Strehl	0.33	0.11	0.44
Pointing (@ peak irradiance Δr) ³	0.0 μrad	3.1 μrad	5.7 μrad
Pointing (@ centroid Δr) ³	0.0 μrad	2.6 μrad	5.7 μrad

¹ Scaled for 372-mm beam² Far field radius of encircled energy for n%. Does not include removal of $\leq 3^{\text{rd}}$ order aberrations.³ Radial pointing i.e., $(\Delta x^2 + \Delta y^2)^{1/2}$

The x and y angular displacement of the far-field centroid for the 50-shot data set is plotted in Figure 2-15. The radial angular displacement is $\Delta r = (\Delta x^2 + \Delta y^2)^{1/2}$. The rms pointing is 2.6 μrad , which fails to meet our requirement of 0.96 μrad . We have identified several possible sources of pointing instabilities in the PAM prototype opto-mechanical design. One of the main sources could be the outlet of a nitrogen purge of the 5-cm rod head located near the optical beamline, which could cause air turbulence. Another possible source is the optical mounts of the regen, which were not designed for vertical placement. These potential problems have been corrected in the first-article design and will be validated by corresponding modifications to the PAM prototype and measurements.




Figure 2-15. Pointing stability (μrad) of far-field beam centroid in x and y coordinate (referenced to a 372-mm beam) is 2.6 μrad . This exceeds the 0.96 μrad requirement because of air turbulence in the PAM, which is being corrected.

The radius of encircled energy in the far field corresponding to 50%, 80%, 90%, and 95% for the 50 shots is plotted in Figure 2-16. (Note the radii are for an uncorrected beam and differ from Figure 2-12). One potential issue we wanted to address is the long-term change in beam quality due to thermal loading in the 5-cm rod over the course of a 50-shot run. As can be seen in the data, there is a small but relatively constant slope indicating there could be additional shot-dependent thermal correction required by the deformable mirror for each shot, which should be predictable in the NIF system model. However, the large fluctuations in the data and the lack of detailed wavefront measurements make it difficult to quantify the rate of change due to thermal loading.

Figures 2-17 and 2-18 are far fields of four subsequent shots (5th, 6th, 7th, and 8th). As can be seen in the figures, the beam breaks up in the far field into several distinct spots for some of the shots. The Strehl changes significantly from 0.17 (shot #6) to 0.63 (shot #8). This performance would be unacceptable for the NIF. We believe that turbulence in the optical beamline due to a nitrogen purge in the 5-cm rod amplifier is also contributing to this problem. We plan to repeat similar measurements after modifications to the prototype have been completed to validate our first-article design.

The large fluctuations seen in the pointing, Strehl, and far-field divergence indicates the PAM opto-mechanical system needs further experimental investigation to determine the source of these instabilities, and then implement and validate corresponding design changes.


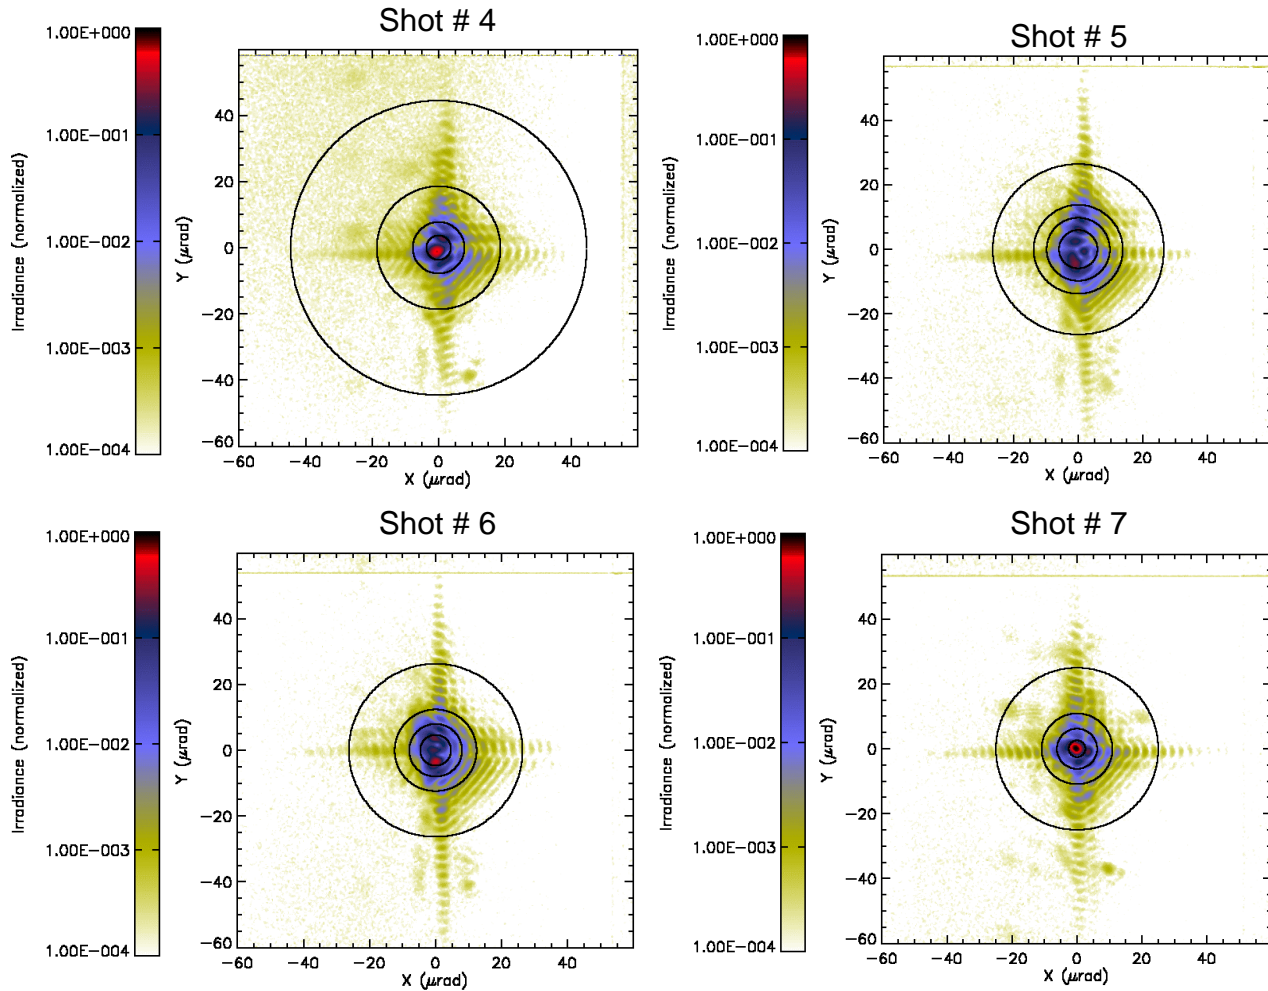


Figure 2-16. Radius (μrad) of encircled energy fractions corresponding (bottom to top) 50%, 80%, 90%, and 95% (referenced to a 372-mm beam).

Figure 2-17. Measured (uncorrected) far-field beam profiles (log scale) for four consecutive 17-J



shots separated by 10 minutes. The circles correspond to 50%, 80%, 90%, and 95% encircled energy. The irradiance scale is identical for all profiles normalized to shot # 8. The large variation in 95% and 90% radius for shot #6 is due to the large background (more detailed experiments are required to determine if this effect is a real part of the beam or an artifact) [referenced to a 372-mm beam].




Figure 2-18. Measured (uncorrected) far-field beam profiles for four consecutive 17-J shots separated by 10 minutes as in Figure 2-17. The irradiance scale (linear) is identical for all profiles normalized to shot # 8 (referenced to a 372-mm beam).

2.8 PAM Procurement Strategy

The first-article PAM will be procured and assembled by a prime integrator using our baseline design on a cost plus fixed fee contract. Two competing contracts will be awarded. The contract will include options for PAM production units. One of the competing prime integrators will be chosen to build the PAM production units.

We have reviewed over thirty companies and have selected eleven as qualified vendors. A prebid conference was held in August in which all eleven companies were in attendance and indicated an interest in building the PAMs. The selection criteria will be based on a balance of four categories, feasibility, applicability, capability, and affordability.

The contract includes the following key tasks to be completed by the prime integrator:

- A project management plan.

- Design changes to the PAM for value engineering (i.e., improve performance, optical hardening, etc.) and reduce manufacturing costs.
- Procurement plan and specifications for fabrication and assembly of PAM components.
- Conformance verification of completed PAMs. The prime integrator is asked to perform verification testing on the completed PAM for confirmation that the PAM subsystems have been assembled correctly. This does not involve energetics testing.
- PAM Characterization Measurements: The prime integrator will be asked to perform tests and record data that will verify performance of the subsystems in the completed PAM as well as output energy, far-field beam quality, and far-field stability.

LLNL will supply diagnostic equipment for conformance and characterization testing. In addition, some components that exist in the PAM assembly, which are difficult to procure or if significant quantities already exist at LLNL, will be supplied as university-furnished equipment. Specifically, we will supply the 42-mm Faraday rotator, the 5-cm rod amplifier assembly, and limited PAM control hardware and software.

2.9 PABTS System Overview

The PABTS system baseline design is based on an open architecture consisting of individual optics mounted on vertical breadboards. The breadboards are mounted on the Pre Amplifier Support Structure (PASS). The breadboards and optical elements are contained in large enclosures that are supplied with high flow rate HEPA filtered air during optics installation. This maintains the required level-100 cleanliness of the optics. The enclosures are backfilled with humid clean air during operation. This maintains cleanliness while reducing turbulence and vibration. Table 2-8 summarizes the PABTS component and integrated PABTS specifications. The PABTS is located after the Input Sensor Package (ISP), which monitors and controls the PAM alignment. The PABTS primary function is to split the PAM beam into 4 beams and provide beamline timing, relay imaging, and energy. The PABTS chain includes: a $55 \times 55 \text{ mm}^2$ Faraday isolator, a 1:4 split, a timing module, and a 6-element vacuum relay telescope (VRT). The timing module is used to synchronize timing between the four separate beamlines. Quad-to-quad timing is synchronized via fiber jumper lengths between the MOR and 48 PAMs. The 6-element telescope is designed to provide accurate beam size adjustment required for the NIF system. In addition, the design also permits both longitudinal and transverse zoom and collimation control. Magnification of the beam size is required to tailor the beam to the amplifier clear aperture and to adjust beam size for beam-to-beam power balance. Relay image adjustments are required to compensate for small changes in path length after downstream optics replacement. Figure 2-19 is the conceptual layout of a PABTS quad. The baseline model of the system is shown in Figure 2-20. In the PASS structure two PABTS quads (1 bundle) share the same set of breadboards and enclosures. Photos of the PASS structure are shown in Figure 2-21.

Table 2-5. PABTS component specifications.

Component	Parameter	Specification
Isolation module	Isolation Operating fluence Clear aperture	> 30dB 5 J/cm ² 45 mm ²
1:4 Split	Beam line energy balance precision	0.1 %
Six-element vacuum relay telescope	Beam size Beam size tolerance Collimation Magnification control Relay distance adjustment	45 × 45 mm ² +0 mm, -0.2 mm 0.02 waves rms 0.1% ± 50 mm
Timing Assembly	Gross timing adjustment (precision) Remote timing adjustment (precision)	± 5 ns (< 30 ps) ± 125 ps (< 1 ps)
Integrated PABTS (does not include contribution of PAM and Input Sensor Package)	Polarization Beam rotation Pointing stability (372-mm beam) Near-field contrast Wavefront (<3 rd -order contribution) Divergence (encircled energy radius) [372 × 372-mm beam] [After .45 waves ≤3 rd -order correction] Pulse timing jitter FM-to-AM conversion	< 0.1 % < 5 mrad rms < 0.35 μrad rms < 4.1% rms < 0.45 wave > 80% in 3.6 μrad < 0.27% outside 33 μrad 4-ps rms 5% rms

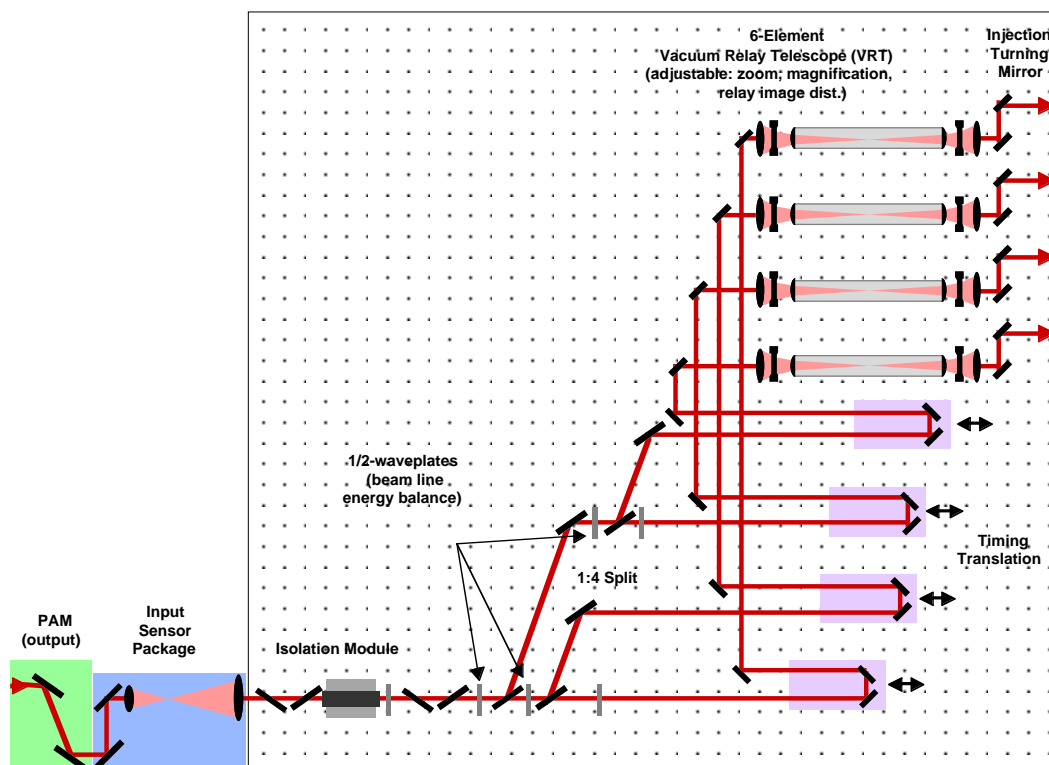


Figure 2-19. PABTS system layout.




Figure 2-20. PABTS conceptual design (full bundle).



Figure 2-21. Photo of a PASS. There are eight PASSes in the NIF. PABTS breadboards for one bundle are mounted on each side of identified cross members.

Very little PABTS prototype work has occurred to date. The large $55 \times 55 \text{ mm}^2$ Faraday rotators from two different vendors have been tested. The 30-dB required extinction of the isolator rotator has been demonstrated (see Figure 2-22).

	EOT	PASSAT
Average rotation angle over aperture	46.8°	46.9°
Minimum extinction	31.1 dB	31.1 dB
Maximum extinction	54.1 dB	51.1 dB
Average extinction over aperture	41.5 dB	41.8 dB
Transmission	98%	96.5%

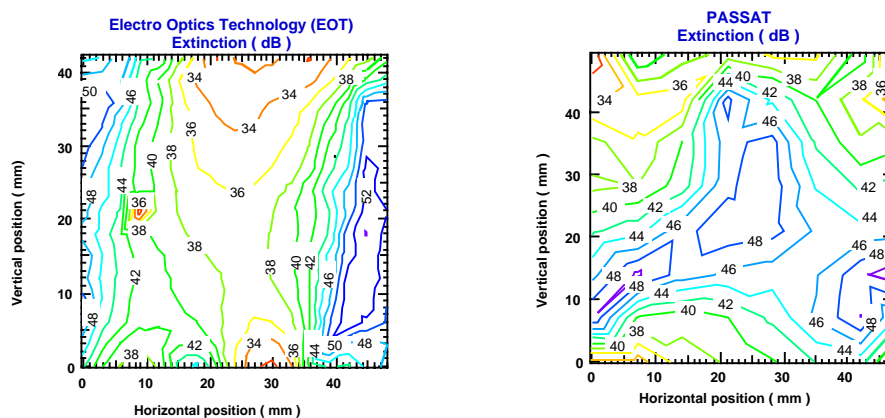


Figure 2-22. Extinction of 55-mm PABTS Faraday rotators supplied by two vendors.

2.10 Procurement Strategy

The breadboards and associated hardware are part of the facility infrastructure and will be provided as government-furnished equipment for the facility subcontractor to install. The PABTS enclosures will consist primarily of sheet metal and standard aluminum channels. After the successful evaluation of one bundle's (8 beams) worth of first articles, these enclosures will be purchased on fixed-price contracts (based on lowest cost) in six-bundle increments, from LLNL drawings. The smaller enclosures, 15 cubic feet or less, will be procured as assemblies. The larger enclosures approach 300 cubic feet and will be assembled in the PABTS corridors, since in the assembled condition they are not transportable in the laser bay. The plan is to purchase the enclosures partially assembled in a standard clean condition in accord with MIL1246C level 300. This meets the cleanliness requirements established by the project. The enclosures will need to be recleaned to level 300 after installation in the laser bay. The HEPA filter system will then be activated.

The PABTS optical mechanical hardware consists primarily of optics (such as mirrors, polarizers, and waveplates) and the associated mounts. The optics and mounts will be procured as assemblies built from LLNL prints. One bundle's worth of optical mechanical hardware will also be evaluated before commitment to procurement of production quantities (see building 381 hibay activities below). The components must be cleaned to Level 100 and assembled and tested in a clean-room environment. This will be a major cost control of this procurement. Assemblies will be delivered to LLNL triple bagged as per existing specifications for precision-cleaned components. The assemblies will be installed on breadboards and inside enclosures supplied with HEPA air during installation.

The vacuum relay telescopes (VRTs) represent the most complex and uncertain of the PABTS procurement packages. Each VRT is 4 meters long with a 3-meter-long vacuum section. It houses six lenses that require the capability for collimation and zoom adjustments. There are seven different VRTs in the NIF that can be adjusted to give 192 unique configurations. The VRT optical design (lens curvature, spacing) will be per LLNL drawings. The requirements of the optical mechanical design (how the lenses are mounted and adjusted) will be provided as specification. The successful VRT supplier will be expected to integrate the LLNL optical design with the optical mechanical specification to provide a unit that meets the performance requirements. One or two vendors will be selected to provide first articles on a fixed-price basis. We will also request bids for production quantities at this time. After evaluation of these first articles by LLNL, we will award the production phase to the vendor having the best combination of design, performance, and price.

The large-aperture Faraday rotators are 20-kg optical assemblies that are required in NIF to isolate the injection laser systems from damaging back-reflected pulses. Both the PAM and PABTS have rotators covered under this procurement. There are 96 such rotators in the NIF. Rotators for the first 12 bundles of the NIF will utilize reworked glass from NOVA. The second 12 bundles will be more costly because new glass will have to be purchased. There are a limited number of suppliers available. Production

capability and competency is an important issue here. The present procurement strategy is to purchase two first articles of the PAM and PABTS varieties with binding bids for production quantities. Contracts will be awarded to one or more suppliers on a best-value basis.

Building 381 Hibay First-Article PABTS/PASS Integrated Testing

Several issues associated with the PABTS benefit from early system integration and testing before installation in the NIF facility. We are currently building a PASS substructure identical to a portion of the PASS planned for the NIF facility in the hibay of Building 381. We will install and test first-article enclosures, optics, and mounts and VRTs in FY00-01. This activity will mitigate several PABTS risk issues.

The PABTS is a complicated and expensive optical mechanical system. The design requires ~100 optical elements per bundle (8 beams). Care must be taken in the design to avoid beam interference with other optics and with the PASS structure that has many cross members. Demonstrating the design early will avoid costly redesign and replacement of procured hardware. Since the current schedule requires approximately one bundle per month installation, significant hardware must be procured before the first bundle is activated and tested in the NIF facility.

The PABTS is located in the PASS, a complex and spatially constrained environment (see Figure 2-16), which will be a difficult work area. In addition, class IIIb and IV laser hazards will be present during installation. Since no safe and cost effective means has been identified to mitigate the laser hazard, the alignment of the PABTS system requires the NIF facility to operate as an active laser facility. This precludes early installation and testing before NIF beneficial occupancy. Building 381 activities will facilitate the development of safe and efficient work procedures needed to maintain the current schedule and eliminate possible accidents.

The PABTS optical elements need to be installed to a level-100 cleanliness requirement. Early installation and cleanliness testing will validate our design and clean installation procedures.

ILS integration requires full system testing. As part of the 381 hibay activities, we will design, build, and test ILS end-of-chain diagnostics package. This package will be designed as a portable diagnostic station and not an in-situ system. This package will be used to acceptance test the first articles in Building 381 and to characterize the output of each ILS beamline during NIF installation. In addition, the package will be used during PABTS maintenance and refurbishment.

The beam-pointing stability for NIF is an important and difficult specification. The PABTS system consists of optics in mounts on vertical breadboards contained in large enclosures. The mechanical stability of the mounted breadboards and the mounted optics design need to be verified, and if problems exist, they need to be corrected. The total beam propagation path in air is 18–20 meters. Beam pointing and break-up due to air density fluctuations could be significant issues. Early

testing would allow for design modifications (e.g., beam tubes, evacuated enclosures, etc.) to be implemented if required. Since these are costly modifications, it does not make sense to incorporate them into the baseline design until problems are identified.

3.0 AMPLIFIER	3-1
3.1 Primary Optical Performance Summary	3-1
3.2 Gain and Gain Profile	3-2
3.2.1 Gain Risk Estimates for NIF—Threat List	3-4
3.2.2 Recommended Mitigation	3-5
3.3 Prompt Pump-Induced Wavefront	3-7
References	3-9
3.4 Amplifier Thermal Recovery (Sutton)	3-10
References	3-16
3.5 NIF Power Conditioning Status Summary	3-17
3.5.1 Summary	3-17
3.5.2 Requirements Summary	3-18
3.5.3 Reliability	3-19
3.5.4 Availability	3-20
3.5.5 Maintenance	3-20
3.5.6 Shrapnel Containment	3-21
3.5.7 Schedule and Overall Strategy	3-21
3.6 Amplab II and First-Article Tests	3-23
First Test of Transporter	3-23
First-Article Slab and Flashlamp Assembly	3-23
Flow Tests of First-Article Slab and Flashlamp Assembly	3-24
Electrical End-to-End Test	3-24
Slab Purge and “Self-Cleaning” Tests	3-24
Disaster Recovery Tests	3-24
3.7 Flashlamp Damage to Laser Amplifier Slabs	3-25
3.7.1 Contamination and Slab Damage	3-25
3.7.2 Description of Flashlamp-Induced Damage	3-26
3.7.3 Contamination and Slab Damage Criteria	3-29
3.7.4 Role of Aerosols in Flashlamp-Induced Damage	3-31
3.7.5 Slab Damage Analysis and Propagation Modeling	3-43
3.7.6 Cumulative Average Damage Distribution for Beamlet	3-47
3.7.7 Logic and Estimated Cost of Slab Refurbishment	3-48
3.7.8 Contaminant Mitigation, Design Strategies, and Operational Strategies	3-49
3.7.9 Retrospective on Nova Amplifier Slabs	3-52
3.7.10 The Small-Scale Flashlamp Damage Testing	3-59

3.0 Amplifier

Ken Jancaitis, Irv Stowers, Chris Marshall, Doug Larson

3.1 Primary Optical Performance Summary

Table 3-1 summarizes the primary optical performance requirements for the NIF amplifiers. It also shows our most recent performance predictions based on experiment and modeling and shows the basis for each prediction. As indicated in the table, our analyses to date show that the amplifiers will meet or exceed 6 of the 8 requirements. The two parameters that have not been verified as being within the required limits are the gain uniformity (peak to average gain ratio) and the prompt wavefront distortion.

Table 3-1 Summary of NIF optical performance criteria and performance predictions.

Parameter	Formal Requirement	Prediction	Basis for Prediction
Average gain coefficient (% cm ⁻¹)	5.0	5.1 ± 0.1	3-D gain model
Gain uniformity (peak/average)	1.05	1.06 ± 0.2	3-D gain model
Mounting stress induced wavefront (waves PV)	0.7	0.3 ± 0.3	AMPLAB experiment
Prompt wavefront distortion (waves PV)	2.7 with < 10% stability	5.5 ± 4 stable to <10%	Prompt pump-induced model calibrated to AMPLAB measurements
Residual thermal wavefront distortion (waves PV) @ 7 hours	2.2 with < 10% stability	0.4 ± 0.25 stable to <10%	Thermal recovery model and AMPLAB measurements
Residual thermal wavefront distortion (waves PV) @ 3 hours	2.2 with < 10% stability	1.5 ± 0.7 stable to <10%	Thermal recovery model and AMPLAB measurements
Prompt pump-induced steering of beam centroid (μrad)	< 8	4	Prompt pump-induced model calibrated to AMPLAB measurements
Pump-induced depolarization (% per slab, aperture averaged)	< 0.05%	<0.01%	Prompt pump-induced model calibrated to AMPLAB measurements

The predicted gain uniformity ratio of 1.06 ± 0.2 is within its error limits of the requirement of 1.05, and propagation calculations show that there is enough margin in the front-end laser system and input beam spatial profiler to make up for this

deficiency.¹ However, the predicted prompt wavefront distortion of 5.5 ± 1 waves is twice as great as the requirement of 2.7 waves. While propagation calculations also indicate that this increase can be accommodated within the limits of the deformable mirror,² the fact that magnitude of the distortion is much larger than what was expected on NIF as well as what was observed in the Beamlet laser system is a cause for concern. For this reason, we are refining our thermal model capabilities to include a more accurate calculation of the wavefront. In the coming year we will use this updated model to reanalyze the Beamlet and AMPLAB optical distortion measurements to resolve this apparent discrepancy.

3.2 Gain and Gain Profile

We used our 3-dimensional ray trace model to predict the gain distribution in the NIF laser slabs. This model has been validated by demonstrating that it accurately predicted the gain variations observed in the AMPLAB prototype amplifier due to differences in aperture position, laser glass composition, and electrical energy delivered to the flashlamps.³

The updated calculations for the NIF amplifier performance took into account the feature changes in the NIF amplifiers that were not present in our earlier amplifiers. These include symmetrically centered side and central flashlamp arrays, thicker blastshields, reformulated LG-770 laser glass, and a polished aluminum (rather than silver) finish on the “minor” reflector elements (slab mask, blastshield mask, corner reflectors, etc.). The characterization of the entire laser chain required 12 separate gain calculations for three distinct positions along the length (interior amplifier, diamond-end amplifier, X-end amplifier), two distinct vertical positions (bottom and next-to-bottom; top and next-to-top are mirror images) and two different glass types (LG-770 and LHG-8). The complete calculation took a week of run-time on our dedicated cluster of UNIX workstations.

A representative gain profile for a diamond-end slab in the bottom aperture using LG-770 glass is shown in Figure 3-1. The gain in the next-to-bottom aperture is slightly higher and shows no significant top-to-bottom roll-off. In this and each of the other gain calculations, all of the optical components were assumed to be pristine with no degradation. The average gain profiles over the entire laser chain for the bottom aperture assuming the baseline 50-50 mix of LG-770 and LHG-8, is presented in Figure 3-2. Under these conditions, the beam-averaged gain is $5.1\% \text{ cm}^{-1}$, and the peak-to-average gain ratio is 1.06.




Figure 3-1 Calculated NIF gain profile for a diamond-end slab, bottom aperture, LG-770 laser glass, at a flashlamp explosion fraction of 20%.




Figure 3-2 Predicted chain-averaged gain profile for the bottom aperture of the NIF amplifiers (50/50 LG-750/LHG-8 mix, 20% explosion fraction).

3.2.1 Gain Risk Estimates for NIF—Threat List

We do not expect all of the optical and electrical components in the NIF amplifier system to go without degradation during the operational lifetime of the laser. With this in mind we have carried out a gain risk analysis for the amplifiers based on both expected vendor performance and component life tests. We used our amplifier gain models to estimate the effects of each component's degradation on the average gain coefficient and stored energy of the laser system. The results of this analysis are presented in Table 3-2, which describes the risks and their effect on the performance of the amplifier. In this table, the gain coefficient and stored energy columns indicate the change in each quantity produced by the degradation of the various components, relative to the baseline gain requirement of $5.0\% \text{ cm}^{-1}$. The risks are separated into either a high- or low-probability designation based on our experience to date. The high-probability risks are those we expect to see without a significant investment in component development or system cleanliness. The low-probability risks are those that we would only expect to see if the quality of the components that are installed in the NIF is that of the worst and not the best pilot-run samples delivered to date.

Table 3-2. Gain risk estimates for NIF.

Risk	Gain coefficient increase/ decrease (% cm⁻¹)	Stored energy increase/ decrease (MJ)	Comments
Baseline model prediction Margin based on latest calculations	+0.1	+0.13	Assuming 50/50 mix of LHG-8/LG750 laser glass
High probability risks			
Blastshield AR coating degradation	-0.13	-0.17	Loss of 4% per surface (based on life test with baseline glass and coating)
Flashlamp reflector degradation	-0.05	-0.06	5% loss of reflectivity (based on first pilot run)
Decrease glass lifetime from water	-0.06	-0.08	4cm ⁻¹ OH absorption at 3 um (baseline is 2 cm ⁻¹)
Thin laser glass from refinishing	-0.11	-0.14	Laser glass thickness reduced from 40.5 to 39 cm (~ 1 repolish)
Power conditioning aging	-0.035	-0.05	Based on FANTM measurements
Low probability risks			
Flashlamp envelope transmission loss	-0.08	-0.10	5% reduction in transmission (based on life test from low cost vendor)
Flashlamp reflector degradation	-0.05	-0.06	Total 10% loss of reflectivity (based on first pilot run)
Thin laser glass from vendor	-0.11	-0.14	Laser glass thickness reduced from 40.5 to 39 cm (~ 1 repolish)
Reoptimize cavity for reduced distortion	-0.07	-0.09	Reoptimized cavity for up to 5x lower prompt wavefront distortion
Fabrication error	-0.03	-0.04	Based on AMPLAB fabrication experience

3.2.2 Recommended Mitigation

Two options have been explored to mitigate the performance loss caused by any of the component degradations described in the preceding section: adding more capacitors to the pulsed power system and/or adding two amplifier modules to the booster amplifier. The performance increases resulting from each of these strategies is shown in Table 3-3. At the same time, it should be noted that each of these strategies to increase the optical performance introduces additional risks to system operation. For example, the addition of more capacitors increases the energy delivered to the flashlamps and decreases their effective lifetime.

Table 3-3. Risk mitigation estimates for the NIF.

Mitigation strategy	Gain coefficient increase/ decrease (% cm-1)	Stored energy increase/ decrease (MJ)	Comments
Add 4 capacitors to pulsed power system	+0.25	+0.32	Flashlamp explosion fraction increased to 23%
Add two amplifier modules to booster amplifier	N/A	+0.81	Increased non-linear effects in booster amplifier

The combined effect of the risk estimates and the mitigation strategies on the total energy stored in the laser system are presented graphically in Figures 3-3 and 3-4. The variation in the total stored energy with flashlamp explosion fraction (number of capacitors ranging from 20 to 24) is plotted as a band of varying density from the most optimistic (no risks) to conservative (all risk realized) risk assumptions for an 11-5 slab laser system and a 11-7 slab laser system, respectively. The graphs show that if only the high-probability risks are realized, the energy storage requirement can be recovered by simply adding the additional capacitors. If, in addition, some of the lower probability risks occur, adding the capacitors will not be enough, and the addition of either the laser slabs or some other mitigation strategy would be required.




Figure 3-3. Variation in the total stored energy with flashlamp explosion fraction (horizontal axis) and for the range of optimistic to conservative risks estimates (colored band) for the baseline 11-5 slab NIF laser system.




Figure 3-4. Variation in the total stored energy with flashlamp explosion fraction (horizontal axis) and for the range of optimistic to conservative risks estimates (colored band) for an 11-7 slab NIF laser system.

3.3 Prompt Pump-Induced Wavefront

The prompt optical distortion predictions for the NIF were obtained by extrapolating from 1-, 2-, and 3-slab measurements and calculations for the AMPLAB prototype amplifier. The measured and model prediction for the 3-slab AMPLAB setup are shown in Figure 3-5, and the corresponding model prediction for the full NIF chain is presented in Figure 3-6.




Figure 3-5. Measured and predicted prompt double-pass optical distortion in the AMPLAB prototype amplifier in the 3-slab-long configuration.




Figure 3-6. Predicted prompt optical distortion for the baseline NIF amplifier chain.

For the sake of comparison, the much smaller measured and predicted optical distortion profiles for the Beamlet laser system are displayed in Figure 3-7.




Figure 3-7. Measured and predicted prompt optical distortion in the Beamlet prototype laser system. The dashed lines indicate the one-sigma measurement uncertainty.

There are several factors that introduce error into the extrapolation of the AMPLAB wavefront measurements to the predicted NIF optical distortions. The AMPLAB distortions include an unknown (and ignored) component coming from flashlamp heating of the return mirror located close to the end of the amplifier. The model for the thermal loading of the glass by the flashlamp light assumes that the temporal and spatial profile of the heat load is identical to that of the pumping of the inversion in the slab. Finally, the magnitude of the heat load was fixed by scaling with an empirically determined constant.

In light of these shortcomings, we have decided to refine the thermal model in the coming year by using the ray-trace technology developed for the gain calculations used to produce an ab initio model for the thermal deposition in the laser slab. The results of this calculation will be used as input to an improved model of the stress and strain in the slab. In this way we hope to resolve the discrepancy between the NIF and Beamlet wavefront distortions.

References

1. W. H. Williams, "NIF Gain with 7/7/99 Gain Calculations," NIF-0036844, August 1999.
2. R. A. Zacharius et. al., "NIF Deformable Mirror (DM) Status and Focal Spot Performance Validation," NIF-0016202, November 1999.
3. A. C. Erlandson et. al., "Physics Basis for the Optical Performance of the NIF Amplifiers," UCRL-ID-132680, January 1999.

3.4 Amplifier Thermal Recovery (Sutton)

One factor that determines the shot rate of the NIF laser system is thermal recovery of the laser amplifiers, which is broken into two categories: (1) wavefront recovery of the laser slabs, and (2) wavefront recovery of the gas columns within or adjacent to the amplifier. Both optical distortions are driven by temperature imbalances, which allows efforts to focus on characterization of the temperature distribution within the laser system. Current understanding of laser slab and gas column wavefront distortion recovery has been previously reported.¹⁻³

As shown in Figure 3-8, current projections of gas-motion wavefront distortion recovery of the amplifiers is strongly dependent on model assumptions and the cooling profile employed. In this figure, projections are given for the cases of ambient temperature flashlamp cooling and for chilled gas cooling, where the latter case considered 1°C chilling of the flashlamp cooling gas for the first 2.5 hours of the recovery cycle. The model applied in this analysis was founded on AMPLAB and Beamlet data.³ For each cooling approach, the performance domain is bounded with a conservative prediction, which included estimates of height scaling differences in the systems and an aggressive prediction that ignored any height scaling differences. These results indicate that ambient gas cooling will likely be sufficient to achieve 7-hour recovery, and that as little as 1°C subcooling of the flashlamp cooling gas will clearly meet baseline recovery requirements, as well as potentially satisfy requirements for an accelerated shot, 3-hour recovery. However, they also serve to point out the uncertainty in our present understanding of gas distortion recovery.




Figure 3-8. Predicted gas motion phase distortions for the NIF system, for both ambient gas cooling and chilled gas cooling. The conservative assumptions incorporate height-scaling effects (see Reference 3), while the aggressive assumptions assume negligible height-scaling differences.

With this uncertainty in mind, we have undertaken an extensive effort to precisely numerically simulate the spatial and temporal character of the gas distortions in the NIF laser chain. The methodology under investigation relies upon precise numerical simulations of the instantaneous flow and temperature fields in the laser amplifier cavities and adjacent beamtubes, which in turn can be related to the optical path length (OPL) variation across the aperture, which is defined as

$$OPL = \int_{\xi} n \cdot d\xi \quad (1)$$

In Eq. (1), n is the index of refraction defined by

$$n = n_o + \frac{dn}{dT}(T - T_o) \quad (2)$$

and ξ is the propagation path through the gas region. In Eq. (2), T is the local instantaneous temperature, n_o is the reference refractive index taken at temperature T_o , and dn/dT is the index temperature coefficient. If we decompose the instantaneous temperature (T) into time-averaged (\bar{T}) and fluctuating (T') quantities, we can in turn decompose the optical path length into time-averaged and fluctuating quantities, i.e.,

$$OPL = \overline{OPL} + OPL' \quad (3)$$

where

$$\overline{OPL} = \int_{\xi} \left(n_o + \frac{dn}{dT}(\bar{T} - T_o) \right) d\xi \quad (4)$$

and


$$OPL' = \int_{\xi} \left(\frac{dn}{dT} T' \right) d\xi \quad (5)$$

The thrust of this effort is to accurately calculate both the spatial and temporal character of the temperature distribution. Once the mean and fluctuating temperature fields are established, the phase distortion through each region can be calculated by applying a ray-trace algorithm. This will produce a phase screen for each discrete region in the amplifier (the triangular cavity regions and the adjacent beamtube regions), which through coherent addition, will account for the appropriate number of laser beam passes.


To calculate the instantaneous temperature field, we have applied the FIDAP computer code,⁴ which is based on the finite-element method and has a generalized meshing capability that is well suited to simulating the buoyantly driven flow in the triangular cavity regions of the amplifier. When calculating a quasi-turbulent flow, a common practice is to apply a low Reynolds number turbulence model that accounts for the velocity and temperature fluctuations through enhanced transport properties. With this approach, however, spatial and temporal fidelity is lost. Since we required a precise understanding of the spatial and temporal structure of the flow, we use the direct numerical simulation approach in which sufficiently small time steps and a sufficiently dense mesh are used, so that the basic transient equations of momentum and energy conservation can be directly applied to calculate the instantaneous character of the quasi-turbulent flow field.

Representative calculated time-averaged (\bar{T}) and fluctuating (T') temperatures in a triangular amplifier cavity are given in Figures 3-9b and 3-9c. These results are for half of one cavity (Figure 3-9a), where symmetry boundary conditions were applied along the plane separating the two halves. Figure 3-9b gives the time-averaged temperature distribution on three slices taken through the cavity and perpendicular to the blastshield surface. The temperature of the slab surface is 0.15°C greater than the blastshield surface temperature, which was at ambient temperature conditions. This temperature imbalance is very near the required minimum based on earlier analysis.²⁻³ As a result of this imbalance, the flow is up the 'hotter' slab surface and down the 'colder' blastshield surface. In these figures, we see a top-to-bottom stratification in the more open regions of the cavity, as well as a pronounced left-to-right temperature variation in the constrained regions near the triangular tips and near the walls.


(3-9a)



(3-9b)



(3-9c)




Figures 3-9a, b, and c. Iso-plots, on three slices taken through the amplifier cavity, of the time-averaged and fluctuating temperature fields. Note that the flow is generally up the hotter slab faces (near side of the figures) and down the colder blastshield faces (far side of the figures): (a) the modeled segment; (b) the time-averaged temperature field; and (c) the fluctuating temperature field.

Figure 3-9c gives the distribution of the instantaneous temperature. Evident in these slices taken through the temperature perturbation field is the presence of a cellular structure that convects with the buoyantly driven flow and at a propagation velocity that is roughly half of the local mean velocity. These cells, which are hotter and colder


than the local time-averaged temperature, represent the index disturbance structure through which light rays will propagate and are the source of the fluctuating optical distortion as defined by Eq. (4). These results demonstrate that the modeling approach is able to capture details of the temperature disturbance structure in this complex three-dimensional geometry.

The temporal history of the temperature at three points in the system are given in Figure 3-10. The points are near the slab face and at an elevation of 84% of the cavity height. This calculation started with a quiescent flow field at a uniform temperature that was the mean of the blastshield and slab surface temperatures. During the initial 200 s of the simulation, the bulk character of the buoyant flow and the periodic disturbances were established. We see in these traces the temporal nature to the fluctuating temperature field that results from cells that are alternately hot and cold, relative to the local time-averaged temperature, propagating past the fixed location in space. There is no indication of any phase correspondence between the traces, which is an apparent artifact of the three-dimensional nature of the geometry and flow.

(3-10a)



(3-10b)



Figures 3-10a and b. Time-history plots at three locations in the amplifier cavity, showing the temporal character of the temperature field: (a) the locations of the three points, relative to the slab and blastshield surface and at an elevation 84% of the cavity height and (b) time-history plots of the instantaneous temperature at the three locations.

A comparison of calculated temperatures with temperatures measured in Beamlet and AMPLAB is given in Figure 3-11. The Beamlet data is from a thermocouple located in the top region of the beamtube adjacent to the main amplifier, the AMPLAB data is from a thermocouple located in the ‘south’ tower, and the calculated value is at location ‘C’ (Figure 3-10a) in a representative amplifier cavity. Because it was difficult to precisely select data sets in which the driving temperature differences were exactly the same as in the calculation, the data sets are offset so that they can be overlayed. It is also important to note that each of the measured temperatures represents an average of roughly 3000 values taken over a 30- to 40-second time frame. This is simply an artifact of the data acquisition system and clearly will damp out any structure with periods less than this averaging time. The point, however, is that the calculated fluctuating magnitude is generally consistent with the magnitudes measured, which provides some validation of our modeling approach.




Figure 3-11. A comparison of calculated temperature disturbances in the NIF amplifier cavity with measurements made in a Beamlet beamtube and an AMPLAB tower. Considering the temporal filtering present in the measured data (associated with the data collection approach), these results show that that temperature fluctuations calculated with the computer model are reasonable.

References

1. *Inertial Confinement Fusion – 1998 ICF Annual Report*, Lawrence Livermore National Laboratory, UCRL-LR-105820-98 (May 1998), pp. 49–50.
2. S. Sutton, A. Erlandson, R. London, K. Manes, C. Marshall, C. Petty, R. Pierce, L. Smith, L. Zapata, J. Beullier, and B. Bicrel, “Thermal Recovery of the NIF Amplifiers,” 3rd Annual Conference on Solid-State Lasers for Application to Inertial Confinement Fusion, Monterey, CA, July 1998.
3. *Physics Basis for Optical Performance of the NIF Amplifiers*, Lawrence Livermore National Laboratory, UCRL-ID-132680 (January 1999), pp. 99–129
4. FIDAP is a finite-element fluid dynamics computer code licensed from Fluent, Inc., Lebanon, NH.

3.5 NIF Power Conditioning Status Summary

Doug Larson

3.5.1 Summary

The 100% Title II design review for the NIF Power Conditioning System (PCS) was held on February 17, 1999. This document presents a summary of the PCS status and outstanding issues and the plans to resolve them prior to installing the NIF hardware.

The key performance parameters required from the PCS appear to be satisfied by the existing design (see Figure 3-12). The delivered energy and corresponding gain coefficient (assuming the baseline amplifier design) meet or exceed requirements. The shot-to-shot pulse reproducibility, an important part of the overall NIF power balance budget, is also well within allowable limits according to measurements made on a first-article module. PCS components that were the source of concern early in the design, including the spark gap switch and the damping elements, appear to be operating reliably.




Figure 3-12. Typical NIF power conditioning module, which supplies up to 2 MJ of energy for 40 flashlamps.

Two areas of concern remain that were raised by the design team and seconded by the 100% review team—reliability, availability, and maintainability (RAM) and shrapnel containment. Both of these issues arise from an examination of the failure data from the first article tests. Extrapolating the first-article experience to predict NIF reliability and availability yields an unsatisfactory result. Similarly, the five catastrophic failures experienced during testing of the first article demonstrate the inadequacy of the module enclosure to contain shrapnel produced by high-energy faults within the module.

A tiger team consisting of explosive containment and ES&H experts as well as PCS personnel was assembled to examine the shrapnel containment problem. Design solutions were identified and verified through analysis. Design changes have been implemented in the design packages. An updated RAM analysis of the system was also completed. Additional data was gathered from a variety of sources on component-level reliability, and it indicates that the design modifications in combination with the use of qualified components will improve the reliability and availability to meet requirements.

3.5.2 Requirements Summary

The status of the design relative to the key requirements is summarized in Table 3-4. The key requirements related to amplifier performance: power, power stability, and lamp-to-lamp uniformity appear to be in good shape, although the absolute power measurements are still being verified for absolute calibration. The areas of primary concern are reliability, availability, and fault containment.

Additional testing is required to verify that design responses to the observed reliability and availability issues are adequate. This testing will be done in the Module Test Stand in Building 391 beginning in July 2000. Shrapnel containment designs have also been implemented.

Table 3-4. PCS requirements status summary.

Key Requirement	Value	Status	Comments
Delivered power pulse—gain	>5%/cm	5%/cm with no margin at 20 capacitors	Approx. 0.1%/cm error bars due to voltage probe issue - in work at SNL.
Pulse shape (power)	360 us at 10% points, nominal	OK	
Shot-to-shot power pulse reproducibility (power balance)	+/-1%	OK	Production power supplies do not exhibit the thermal effects observed on the SNL prototype.
Lamp-to-lamp power pulse uniformity	+/-3%	OK	Tests at SNL show PCS is within spec. Lamps use nearly all uniformity budget. See NIF 0018787 for sensitivity analysis.
System reliability	92%	95%	Bottom-up reliability assessment complete, but needs verification through testing.
System availability	4 hours per year of unplanned maintenance	.9983	Requirement assumes all failures diagnosed and corrected between shots—some faults may take longer than allowed 7 hours for current design.
Containment of shrapnel from faults	Contain within capacitor modules	OK	Enclosure redesigned according to tiger team recommendations.
Grounding for personnel safety	< 500-V touch potential	TBD	A model validated with sub-scale experiments indicates the design works, but needs NIF-scale test.

3.5.3 Reliability

The overall system reliability and corresponding implications for the NIF's reliability are a concern of the design and review teams. The overall reliability status at the component or subassembly level is summarized in Table 3-5. Although the data to date using prototype parts suggests reliability as low as 75% for the NIF system, the new analysis indicated a reliability as high as 95%. Nearly all of the unreliability in the prototype results came from the capacitors, switch system, and preionization. The capacitors had a design flaw identified prior to starting NIF qualification testing. The qualification data demonstrates that the NIF capacitors will be substantially better. The switch system had numerous prefires and misfires. Analysis of the data suggest that problems with the gas system that sets the air pressure in the spark gap is the cause of most or all of the problems with the switch. The gas system has been redesigned to address the issues. Design modifications have been incorporated into the preionization system to resolve the problems observed. Other failures during prototype testing included charging supplies and dump resistors. The circuit that protects the charging supply from transients during normal and fault events was deemed inadequate during prototype tests. The circuit has been redesigned and tested to improve the charger

reliability. The dump resistor has been redesigned to address the failure mode identified.

The fourth column in Table 3.5 shows the reliability for each of the key components. Achieving these goals in all cases would result in a margin of 6x in PCS reliability. The goals, while aggressive, are reasonable, and plans are in place to reassess the reliability status of each component.

Table 3-5. Overall PCS reliability status at the component or subassembly level.

Subassembly/ component	Allocation (Failed shots per year out of 770 baseline)	Bad shots per year due to (worst case):	Bad shots per year due to (best guess/goal):	Comments and planned actions
Capacitors	31	100	5	Complete qualification tests; use best value source selection
Charging supply	5	20	1	Procure early and measure reliability; test between shots
Switch/trigger system	12	50	4	Gas system is major (possibly only) problem; new design addresses problem and is in test; define preventative maintenance procedures
Preionization system	8	10	0	Shouldn't ruin shot even if it fails; reliability being measured
Cables and connections	2	5	0	Qualify vendors and acceptance test cable
Controls	3	2	0	No problems anticipated
Other	1	5	0	No problems anticipated
Totals	62	192	10	

3.5.4 Availability

The existing unavailability allocation for the PCS is only four hours per year and is based on the assumption that all unplanned maintenance can be completed between NIF shots on a noninterfering basis with operations. While this assumption is largely true, experience with Nova and common sense imply that the actual integrated annual delay attributed to the PCS will be substantially greater than four hours. An ECR is being written to increase the allocation to 4 hours per week, a more realistic goal. The prototyping effort identified maintenance activities likely to exceed the available time between shots, and the design and maintenance strategies are being modified to speed up certain activities such as replacing a failed capacitor.

3.5.5 Maintenance

The only components in the system that require routine planned maintenance are the main switch and the preionization switch. The electrode lifetime of the main switch

is limited to roughly 2000 shots. A spark plug in each spark gap enhances the triggering simultaneity of the devices and must be replaced every 500 shots due to erosion. In addition, the gas lines and filters that feed the main switch get clogged due to electrode debris and must be replaced when the switches are refurbished. The preionization switch lifetime is approximately 4000 shots. Erosion of the switch electrode eventually leads to excessive jitter in the switch. The trigger electrode must be replaced to maintain reliable and predictable operation. The switch “wear-out” for both switches is very predictable, and maintenance can be scheduled to avoid impacting system reliability or availability.

3.5.6 Shrapnel Containment

A tiger team of explosive containment, ES&H experts, and PCS personnel was assembled to examine the shrapnel containment problem and to develop design concepts for mitigating the shrapnel problem. This team analyzed the problem in great detail using photographic and forensic evidence from the catastrophic failures that were experienced on the prototype module at Sandia. This evidence was used to estimate the quantity and impact of the energy released during these failures. A “louvered door” design for the module enclosure was proposed that would relieve the estimated pressures generated within the module and contain 100% of the shrapnel generated by the failed components. This proposed concept has been incorporated into the design of the module. Analysis of the door design and the enclosure predicts that the present design will alleviate the shrapnel problem. Since it is difficult to design and implement a conclusive test of this design, a philosophy was adopted that provides two “layers” of protection. A decision was made to shield the walls of the capacitor bays in addition to the design changes of the module enclosure. The capacitor bay shielding is a minimum of 1 and 1/8-inch plywood. This thickness of plywood was chosen because estimates indicate that it would stop our worst-case observed shrapnel. This design change to the capacitor bays is currently being implemented.

3.5.7 Schedule and Overall Strategy

The summary schedule of risk mitigation and first-article testing is shown in Figure 3-13. Capacitor vendor qualification has been completed, and contracts have been awarded to the two vendors with the most reliable capacitor designs and the lowest cost. A new prototype module with the final design configuration is being built for testing beginning in late July 2000. This activity will include reliability testing and integrated system testing. The integrated system test will include both the PCS and the amplifier FAU hardware. This testing will occur in parallel with the procurement, fabrication, and delivery of initial production modules from two vendors. Multi-module tests will be performed in mid to late FY01 using the initial production modules. The PCS RAM analysis will be updated to incorporate current reliability data and will be used to predict the overall subsystem reliability.

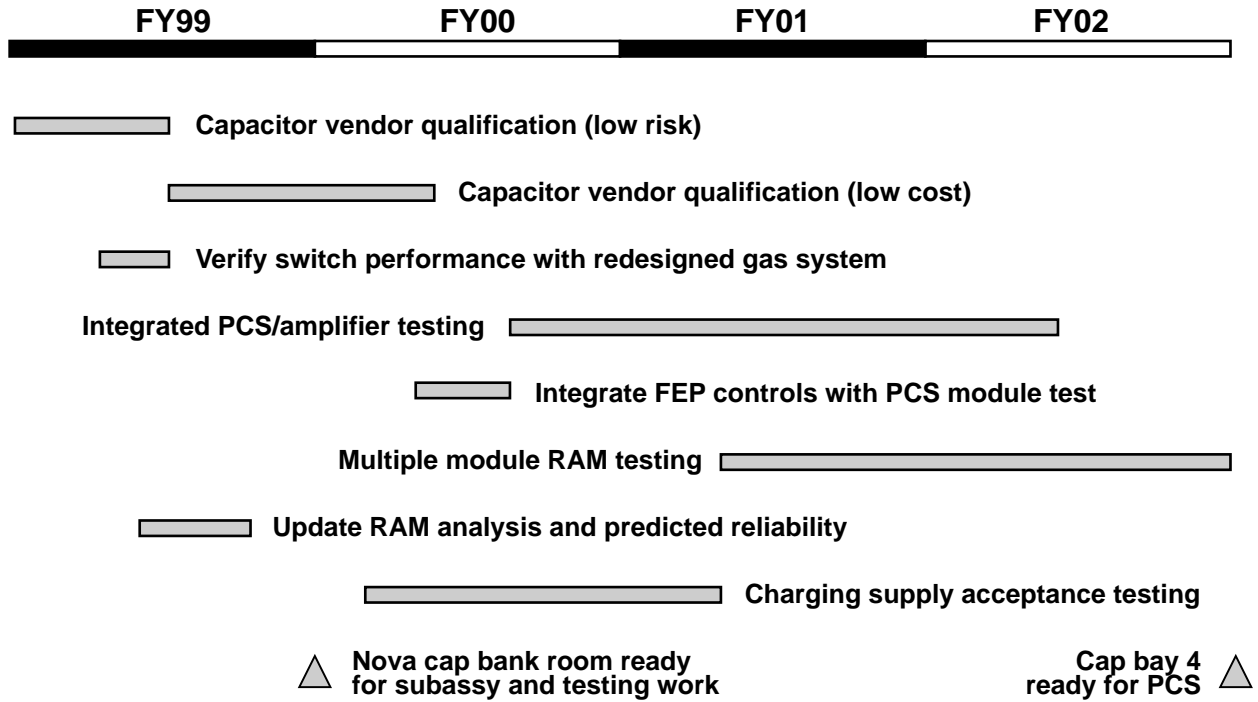


Figure 3-13. Summary schedule of risk mitigation and first-article testing.

3.6 Amplab II and First-Article Tests

The primary purpose of Amplab II is to perform a full, integrated system test of the NIF Article I hardware. The Amplab II experiments will not include scientific experiments such as gain measurements, wavefront measurements, nor thermal recovery measurements. Individual components and subassemblies can be, and are, tested under specific, controlled conditions, but system integration tests often reveal unforeseen problems. There has been much discussion about the difficulties and potential cleanliness problems with the NIF amplifiers. The Amplab II experiments will address a number of these issues, particularly, how clean are the amplifiers, how dirty do they become, and what is the resultant damage from the contamination? A good deal of information was learned from Beamlet and Amplab I. The Beamlet experiments showed cleanliness was of the utmost importance, but if the proper procedures and protocols were followed, cleanliness could be maintained at an acceptable level. The Amplab II experiments will allow us to determine what level of cleanliness is necessary for the successful and damage-free operation of the amplifiers. If the Amplab II experiments are not performed, extra precautions will have to be taken to guarantee a cleanliness level that may, in fact, be more than is required. In terms of cleanliness, it becomes a question of funding Amplab II to determine the requirements, or funding extra-cautionary cleanliness procedures and protocols during NIF construction.

The experiments to be performed are listed and outlined below.

First Test of Transporter

A clean and assembled FAU will be loaded onto the transporter and driven down the street to the OAB and returned. Aerosol measurements will be taken inside the canister and inside the transporter during transportation. Cleanliness of the various surfaces will also be measured both prior to departure from, and after return to, Bldg. 391. Prior to installation, all of the slabs will be measured for obscurations using the flatbed scanner. After the return to Bldg. 391, some of the slabs will be removed and rechecked for cleanliness using the flatbed scanner, and these data will be compared to those of prior to installation. If the cleanliness level is not maintained to a reasonable level, an investigation will take place to determine the source of the contamination and to remove the source.

First-Article Slab and Flashlamp Assembly

The first-article slab cassettes and flashlamp cassettes are to be integrated in the FAU and LRU to check for form, fit, and function. The amplifier will be between 2 and 4 slabs long. Any assembly errors or difficulties will be noted and, if possible, corrected. The cassettes will be inserted and removed some number of times and checked for cleanliness level as well as for generated aerosol concentration. Slab obscurations will be measured using the flatbed scanner. These measurements will be made before

insertion and again after removal. Aerosol concentration will be measured inside the FAU.

Flow Tests of First-Article Slab and Flashlamp Assembly

The gas flow connections to the assembled Article I hardware will be tested for form, fit, and function. Any assembly errors or difficulties will be noted and, if possible, corrected. Flow will be measured in both the central and side flashlamp cassettes to check for flow rate and central-to-side balance. Flow will also be measured in the slab cassettes to determine flow rate and flow balance throughout the amplifier. Any deviations from expected values will be noted and, if possible, corrected.

Electrical End-to-End Test

The pulse-power system will be tested with real flashlamp cassettes using Article I hardware. The parameters to be examined are voltage and current waveforms along with cassette-to-cassette uniformity. The form, fit, and function of all electrical and pulse-power interconnects will be tested and verified.

Slab Purge and “Self-Cleaning” Tests

Using the fully assembled Article I hardware, aerosol and purge-rate measurements will be made prior to, during, and after each shot. Aerosols will be measured to check for removal rate, and flow conditions will again be monitored for uniformity. Using well-documented slabs (as measured by the flatbed scanner), the slab cleanliness will be monitored as a function of flow rate and flashlamp explosion fraction. We hope to gain an understanding of the optimum flow rate to minimize flashlamp damage. We also hope to determine whether or not it is useful and beneficial to slowly increase the flashlamp energy to “self-clean” the slabs. If it is useful, we will establish a “start-up” procedure for NIF amplifiers.

Disaster Recovery Tests

Should any disastrous or “off-normal” events occur during Amplab II testing, we will have an opportunity to investigate in-situ cleaning. Whatever occurs, the cleanup procedure will be well documented, including photographs, before and after cleanliness level measurements, and before and after aerosol measurements. If there are no disastrous events, a disaster will be simulated. All problems and difficulties will be noted and, if possible, corrected.

Disaster Recovery Tests

Amplab II gives us an opportunity to do system integration tests on actual Article I NIF hardware. System integration tests often illuminate unanticipated problems and difficulties. The Amplab II experiments allow us to catch and correct these problems on actual NIF hardware in a more manageable and forgiving laboratory-type environment.

3.7 Flashlamp Damage to Laser Amplifier Slabs

Irving Stowers, John Honig, Steve Letts, Mary Spaeth, Ken Manes, Clay Widmayer, Chris Marshall, Joseph Menapace, Mike Feit

3.7.1 Contamination and Slab Damage

3.7.1.1 Introduction

For many years it has been observed that various levels of damage occur to laser amplifier slabs during normal laser operations. The damage phenomenon has nearly always been associated with the introduction of extrinsic (external) particulate contamination such as the inadvertent exposure of gaskets to laser and/or flashlamp light resulting in the generation of massive amounts of burned elastomer particles which form an aerosol and eventually settle everywhere inside the laser cavity. A somewhat related intrinsic (internal) damage has been associated with the inclusion of small platinum particles within the laser glass. These are opaque to the laser light, heat rapidly under flashlamp and laser light, and literally explode within the laser glass. This intrinsic damage problem has been eliminated by changes to the glass fabrication process. Elimination of the extrinsic damage problem will require (1) a more fundamental understanding of the damage generation mechanism, (2) a more precise understanding of the types of contaminants that lead to slab damage, (3) identification of all source of the aerosol particles, (4) development of cleaning techniques to provide sufficient initial cleanliness, and (5) the selection of low outgassing and high damage threshold structural materials that will minimize the formation of aerosols in the presence of laser and flashlamp light.

3.7.1.2 Laser Cavity Environmental Conditions

English¹ has estimated the range of expected scattered laser light in the NIF laser cavity and expects that it will not exceed 50 mJ/cm². Unfortunately, the laser damage threshold of most organic materials is similar to this upper limit of the expected scattered light level within the laser cavity, and therefore, the use of organic coatings (paint) and exposed elastomer gaskets has been prohibited. For comparison, laser fluence in the output laser pulse is expected to reach a peak of 18 to 20 J/cm² and last three to twenty nanoseconds, and flashlamp fluence in all amplifiers is expected to be 10 J/cm² for a 350-μs duration flashlamp pulse. At these fluence levels, nearly all elastomers would sustain extensive damage and, therefore, will need metallic shine shields.

¹ Edward English, *Estimate of General Stray Light Levels in Main Laser Cavity*, NIF-0004093, Dec. 5, 1997.

3.7.2 Description of Flashlamp-Induced Damage

3.7.2.1 Observations of Slab Damage

Maintenance operations on Argus, Shiva, Nova, and Beamlet have observed various levels of damage to the laser amplifier slabs. As early as Argus, it was also observed that this damage occurred without the need for laser propagation; that is, damage was observed after only flashlamp exposure. Figure 3-14 shows a highly magnified image of a slab damage site induced from a single flashlamp firing.




Figure 3-14. Microscopic image of laser slab damage resulting from purposeful contamination of borosilicate glass by low-density polyethylene and then exposure to a single flashlamp firing. The absorbing contaminant particles are rapidly heated by the flashlamp light to the plasma temperature of 10,000°K.

Usually, laser-induced damage initiated by absorbing particles embedded in laser glass surfaces results in small pits whose internal morphology resembles a geode. Particles resting on or near the glass surface, on the other hand, become small UV-emitting fireballs as they come into radiative equilibrium with the ~ 1 -ev flashlamp black-body temperature. Because of the high thermal expansion coefficient and low tensile strength of the laser glass, it crazes, and shallow surface cracks quickly form on the surface. The mechanical properties of laser glass are such that glasses with a high thermal expansion coefficient, low thermal conductivity, and low tensile and fracture strength will be most susceptible to such flashburn crazing (fracturing) when exposed to intense flashlamp light. Table 3-6 lists the thermal shock parameter (TSP) for a variety of laser glasses and crystals. The TSP is a nondimensional parameter that provides an indication of the relative sensitivity of material to crazing. Lower numerical

values indicate an increased sensitivity to thermally induced crazing, and the two phosphate glasses, LG770 and LHG-8, being produced for the NIF, are particularly sensitive to this form of fracturing.

Table 3-6. TSP for several laser glasses and crystals. The TSP is a measure of the sensitivity of a brittle material to crazing fracture.

Glass / Crystal Type		Thermal Shock Parameter ²
Glasses		
	Zerodur silicate	10,678
	SiO ₂ fused quartz	1,946
	ED-2S silicate	538
	LHG-5S phosphate	195
	ED-2 silicate	123
	APG-1	110
	LHG-5 phosphate	92
	LHG-8	68
	LG-770	64
	LG-750	46
Single Crystals		
	Al ₂ O ₃ sapphire	9,031
	MgO periclase	4,971
	BeAl ₂ O ₄ alenandrite	3,812
	SiO ₂ quartz	1,961
	Y ₂ O ₃ yttria	991
	YAG garnet-oxide	788
	MgF ₂ scheelite	585
	CaF ₂ fluorite	338
	ZrO ₂ zirconia	117
	KZnF ₃ perovskite	98

3.7.2.2 Demonstration that Flashlamp Irradiation of Slab Surfaces Induces Damage

Menapace³ has demonstrated that slab damage can be initiated by any of the contaminants listed in Table 3-7. An opaque particle of radius, a , will intercept and be heated by flashlamp fluence of 10 J/cm². Particles with no effective thermal sink must rise in temperature until they are able to radiate their energy. The fireball so produced

² (1-Poisson's ratio) x (thermal conductivity / thermal expansion coefficient) x (fracture toughness / elastic modulus / (crack length)^{0.5})

³ Joseph Menapace; Slab damage induced by seeding laser glass with various contaminants and exposing them to limited flashlamp exposure.

illuminates the nearby glass surface with a spectrum, including UV wavelengths to which the glass is highly absorbing. At some radius, A , away from the fireball on the glass surface, the deposited UV energy falls below the level sufficient to exceed the critical stress for fracturing the material. Since the lamp power coupling to the opaque particle, as well as its reemission, must scale as the particle area, and the surface burn area also scales linearly with reemission power, the flash-burn radius should scale linearly with particle radius; i.e., $A \propto a$. Consequently, the rapid heating of an opaque contaminant results in local heating of the glass surface, which can be relieved by the large-scale but shallow fracturing of the glass. Shallow buried laser-heated particles, by contrast, lead to pits approximately one-third as deep as they are across, where the volume of material removed scales with the absorbed fluence. The most profound result of this analysis is that the resulting size of the fractured area scales linearly with, and can be significantly larger than, the size of the contaminant that initiated the damage. Menapace found that the size of the damage site exceeded the size of the contaminant by a factor of 7.7 as shown in Figure 3-15.

Table 3-7. List of contaminating materials that have been used to demonstrate that any opaque particle placed onto laser amplifier glass and other glasses with similar thermomechanical parameters will crack and craze when exposed to as little as a single flashlamp firing.

Arizona road dust
Carbon black
Clean-room garment fibers
Clean-room wipers
Kapton
Lens tissues
Plexiglas
Polyethylene
Polystyrene spheres
Sherwin-Williams Polane paint
Skin flakes
Teflon
Voranol (W/ 50% ATH)




Figure 3-15. Results from experiments involving the purposeful contamination of laser glass surfaces and the subsequent irradiation of the surfaces with a one- to three-flashlamp exposure. The damage after irradiation is typically 7.7 times larger than that caused by the initial contaminant.

3.7.3 Contamination and Slab Damage Criteria

3.7.3.1 Scattered Light and the Maximum Acceptable Damage Size

A criteria is necessary to establish an acceptable level of laser slab damage. The criterion can be used to compare previous damage observations and to establish slab refurbishment rates. Hunt⁴ has suggested that the scatter loss from all sources (scratches, digs, and contamination) should not exceed 0.1–0.2% during one round trip from one cavity pinhole to another. Assuming 11 slabs in the cavity amplifier (which are double-passed) leads to an initial scatter loss of 2.5×10^{-5} /surface. It is reasonable to allow this scatter loss to increase by an order of magnitude during normal laser operation (to 2.5×10^{-4} /surface or 1–2% for the 11 slabs in the cavity amplifier section) before laser slabs need to be refurbished.

Hunt and Widmayer⁵ have modeled beam modulation induced by opaque contaminants (or damage sites) of various sizes and at various locations on slab surfaces. The modeling predicts the beam modulation on spatial filter lenses and

⁴ John Hunt, private communication concerning acceptable scatter losses in laser systems.

⁵ Clay Widmayer and Paul Renard, "Survey of damage threat to vacuum barriers by defects and inclusions in NIF optical components," NIF-008158, March 15, 1998.

reports its intensity relative to the damage threshold of the spatial filter lens. They have concluded that opaque contaminants or damage sites up to 500 μm diameter on amplifier slabs pose no threat to downstream spatial filter lenses, and they further conclude that the detection of a single opaque contaminant or damage site on a slab surface of 2000 μm (2 mm) diameter represents a safe removal and refurbishment criteria.⁶

3.7.3.2 Optical and Structural Cleanliness Requirements

Initial optical and structural surface cleanliness levels are derived from the calculated scatter loss resulting from the size distribution of various cleanliness levels. Cleanliness levels are comprehensively defined in MIL-STD-1246C⁷ and are shown diagrammatically in Figure 3-16.




Figure 3-16. MIL-STD-1246C defines a family of particle contaminant levels named for the maximum particle size expected on each square foot of surface area. A Level-100 surface should have no more than a single 100- μm particle on each square foot but may also have up to 1800 particles per square foot $\geq 5 \mu\text{m}$.

The area obscured by any given size distribution (or Level of cleanliness) as defined in MIL-STD-1246C is easily computed. For example, a Level-100 optical surface will obscure (or scatter out of the beamline) only 1×10^{-6} of the incident light compared to

⁶ Clay Widnayer, Ken Manes, John Hunt, and Jack Campbell, "Safety of NIF vacuum barriers from failure due to laser-induced damage," NIF-0034740, September 7, 1999.

⁷ MIL-STD-1246C, "Product Cleanliness Levels and Contamination Control Program," Institute of Environmental Sciences and Technology, Mt. Prospect, IL.

the initial scatter loss criteria of 2.5×10^{-5} mentioned in Section 3.7.3.1. Therefore, optical surfaces could be allowed to become contaminated (or damaged) to roughly Level 180 before the distribution of contaminants would result in a scatter loss equal to the criteria of 2.5×10^{-5} mentioned in Section 3.7.3.1. However, it is impractical to allow this level of contaminants to accumulate as Level 180 will statistically allow a single 234- μm -diameter particle per slab surface and furthermore, each 234- μm particle may result in a 1800- μm -diameter damage site ($234 \mu\text{m} \times 7.7$ as discussed in Section 3.7.2.2). This damage size closely approaches the maximum size allowable to prevent possible damage to spatial filter lenses. Therefore, both initial and operational optical cleanliness levels must be maintained at least one order of magnitude cleaner than defined by Level 202, which statistically leads to a single 2-mm damage site per slab surface. Level 202 divided by 10 yields a maximum optical surface cleanliness of Level 114.

3.7.4 Role of Aerosols in Flashlamp-Induced Damage

3.7.4.1 Possible Origin of the Contamination that Lead to Flashlamp-Induced Damage

A survey of likely sources of contaminants that could find their way onto optical surfaces includes the items shown in Table 3-8. Data on the actual rate of slab damage as a function of the number of flashlamp shots is rare, however, there is growing evidence that the damage occurs early and seems to reach an asymptotic level after a few hundred to a few thousand shots. This corresponds to several of the possible contaminant sources as described in Table 3-8.

Table 3-8. Possible sources of amplifier contaminants. The appearance of these contaminants and the corresponding optical slab damage may be related to the availability of the contaminants inside the amplifier.

Possible source of contaminants	Possible dependence on number of shots
Contaminants left after manufacturing and optical polishing	Likely to occur very early, possibly growing with laser or flashlamp irradiation
Contaminants introduced by installation or removal activities	Likely to occur just after installation or removal of nearby slab cassette
Impurities caught in surface scratches	Likely to occur very early, possibly growing with laser or flashlamp irradiation
Aerosols generated from NVRs	Likely to depend on thickness of nonvolatile residue (NVR) film. A first guess might be 150–300 shots to remove a monolayer of NVR
Aerosol generated from elastomers	Likely to be more linear with the number of flashlamp shots and to continue until the elastomer is entirely consumed
Sol gel coming from the blastshield	Likely to occur very early, possibly growing with laser or flashlamp irradiation
Vibration between slab mask and slab surface	Likely to be linear with the number of flashlamp firings and concentrated near the edge.
Evanescent coupling between the slab and the slab mask at contact points	Likely to saturate after a few shots as contact point “burn off”

3.7.4.2 Aerosols as a possible source of slab damage

Stowers has reported the creation of an aerosol in laser amplifiers immediately after the firing of the amplifier. First observed on the Shiva laser circa 1980, it was reconfirmed on Beamlet in 1997. The background aerosol level in an unfired nitrogen flushed amplifier is Class 1–10 (1–10 particles/ft³ > 0.5 µm). Immediately after firing, this aerosol level rapidly jumps to Class 100,000 to 1,000,000. These aerosol levels are not sufficiently dense to be easily seen with the unaided eye, and furthermore, the typical NIF office will have a aerosol around Class 10,000, whereas the typical LLNL outdoor aerosol will be Class 100,000.

Aerosols that are not flushed away relatively quickly will linger within an amplifier until the particles settle out of the air at a rate based upon the particles' size and density. Particles with a density of 1 g/cm³ and a diameter of 5 µm will settle at the rate of only 0.15 ft/min. To reduce the chance of these particles landing on laser slabs, the slabs are mounted on edge so that the minimum surface area faces upward. However, electrostatic charges induced by the laser beam can significantly counteract the forces of gravity and dramatically increase the particle deposition rate onto vertical surfaces.

Damage to laser slabs occurs subsequent to shots that generate aerosols when settled aerosol particles absorb flashlamp light and rapidly heat. The heated particles thermally stress the laser glass and result in local cracking or crazing as discussed in Section 3.7.2.1.

The aerosols that have been observed in amplifiers are believed to be generated when flashlamp light irradiates organic films, a.k.a. nonvolatile residue (NVR) on structural surfaces. The flashlight light fluence of 10 J/cm² in 370 µs is sufficient to pyrolyze and decompose many organic films on structural surfaces.

A very clean structural surface with only a few monolayers of NVR will have 0.1 mg/ft² and a NIF amplifier with 56 ft² of internal surface area will hold nearly 6 mg of NVR. When compared to a typical aerosol of Class 100,000, which only represents 0.5 µg/ft³ (12 µg/amplifier) it is apparent that even very clean structural surfaces can be the source for many flashlamp firings. If the amplifier contains any exposed elastomers, they could represent a nearly indefinite source of aerosols.

Manes has considered that the aerosol that causes damage does not have to settle onto the glass surface, it only needs to be within a few diameters of the surface when the flashlamps are fired. The aerosol will rapidly heat to nearly the temperature of the plasma within the flashlamps (10,000°K), and the hard UV radiation emanating from the hot particle can result in radiation damage to the glass as described in Section 3.7.2.2.

3.7.4.3 Aerosol Particles Collected during AMPLAB Experiments

Experiments in the B391 AMPLAB facility were designed to demonstrate clean amplifier construction and operations. Two amplifier modules were precision cleaned

and assembled with a small construction difference. One unit was equipped with standard elastomer gaskets, while the other was constructed without any elastomers, to determine the extent to which the elastomers contribute to aerosol production and subsequent slab damage.

The aerosol generation was measured utilizing a multiport airborne particle counter⁸ with up to 8 sampling ports located at various elevations within the four-slab-high amplifier. Two additional ports were used to sample the room aerosol and the airflow through a high-efficiency filter to provide a near-zero aerosol measurement during each cycle through the ten sampling channels. The typical aerosol measurement observed immediately after a flashlamp firing was initially Class 100,000 (100,000 particles/ft³ $\geq 0.5 \mu\text{m}$) at any sampling port. The invariance in particle concentration with location is an indication of the significant thermally induced circulation established by radiant heating of the blastshields within the amplifier. Figure 3-17 shows a typical one-channel multishot aerosol measurement taken in the amplifier without elastomers. The peak aerosol concentration slowly decreased, indicative of the slow removal of NVR contaminants from the walls of the amplifier. The multiport particle counter caused a typical nitrogen flow of 4 ft³/min through the amplifier, which accounts for the exponential decrease in particle concentration that occurs after each shot and is particularly noticeable after the final shot.




Figure 3-17. Aerosol measurement showing an initial peak concentration of Class 100,000, which slowly decreases with each flashlamp firing due to slow removal of organic contaminants from the walls of the amplifier.

⁸ Climet 10-port particle counter with 1-cfm flow rate per channel.

3.7.4.4 *Particles from AMPLAB were Examined for Morphology and Composition*

The composition and morphology of the aerosol was ascertained by sampling the aerosol with a cascade impactor, which forces a rapidly moving stream of aerosol to impinge onto a collection surface. The sample surface was photographed in a scanning electron microscope and is shown in Figure 3-18. A higher magnification image shown in Figure 3-19 shows that the particles that are easily counted by the multiport particle counter are in fact composed of many 10–20 nm nodules of nearly pure carbon. It is believed that the smaller particles form during the short time period that the flashlamps illuminate the structural walls of the amplifier, whereas the 3- μm particle clusters form due to the rapid collision of these smaller particles and form shortly after the flashlamps extinguish, and the heated gas rapidly mixes the smaller aerosol particles. Because the smaller nodules do not pack densely, the resulting clusters have a density of only 1/10 bulk density. This results in a significantly lower settling velocity and a much longer aerosol lifetime than would be expected with fully dense particles, which would settle out of the surrounding gas much more rapidly than do these low-density particle clusters.




Figure 3-18. Low-magnification scanning electron microscopy (SEM) photograph of aerosol particles extracted from amplifiers tested in AMPLAB. The total number of particles collected in the cascade impactor is estimated to be 6×10^6 , which represents only 25 ng per shot.



Figure 3-19. High-magnification SEM photograph of aerosol particles from the AMPLAB amplifier. The typically 3- μm particles are composed of 20–30 nm nodules that are made of nearly pure carbon, indicative of decomposed organic matter.

3.7.4.5 What Do We Know about Aerosols and Damage?

Table 3-9 lists a series of concepts that have been assembled to assist in understanding the aerosol damage phenomenon and our level of confidence concerning the observation or concept. At this time, we have a consistent explanation of the cause of laser slab damage, but we have not yet demonstrated the elimination of all slab damage by preventing aerosol formation. This experimental work is being planned for FY2000. One of the experimental sequences will be to distinguish between the coincidental existence of aerosols and slab damage and to clearly demonstrate a cause-and-effect relationship between damage and aerosols.

Table 3-9. Aerosol and damage concept versus our confidence in the concept.

Aerosol/damage concept	Level of confidence
Aerosols are present just after a shot in all ICF laser cavities	Measured in Shiva, Nova, Beamlet and AMPLAB
NVR that are subjected to flashlamp light produce aerosols	Measurements on AMPLAB and in the small-scale facility are very repeatable
Elastomers (polymers) that are subjected to flashlamp light produce aerosols	Measured in small-scale tests Suspected from full-scale results
The aerosol “class” produced by an NVR or elastomer is a function of the 4 th power of the flashlamp fluence	Measured during small-scale tests. Measured in AMPLAB in two ways LRU loaded vs. unloaded w/ glass PILC vs. full power lamp firing
The aerosol size distribution found after a flashlamp shot roughly obeys a power law (Figure 3-17)	Observed for all aerosol measurements
The size distribution of slab damage roughly follows a power law	Found on Shiva, Beamlet, and AMPLAB
The total amount of damage increases as a function of the number of shots	Measured, but the relationship is not well understood
Laser glass damages under flashlamp illumination if particles are on the surface	Measured after seeding surface with contaminants
Observation of “plasmoids” in the amplifier beamline during flashlamp firing	Historically observed
Phosphate glass has poor thermal shock resistance	Measured by others

All measured aerosol particle size distributions roughly follow a power law as shown in Figure 3-20. The slope of the cumulative size distribution is roughly -2.2 , which is similar to the slope of the aerosols defined in Federal Standard 209 used to define the classification of clean rooms. By implication, the same mechanisms that form naturally occurring aerosol are manifesting themselves in the aerosols observed in our laser amplifiers, and they are able to form a “standard” aerosol within one minute of a flashlamp firing.




Figure 3-20. Cumulative size distribution of aerosol generated after typical flashlamp shot. The flashlamp-generated aerosol is plotted along with the aerosol size distribution used to define naturally occurring aerosols in clean rooms as defined by Federal Standard 209. The data for the aerosol labelled 8/26/98 was taken from Figure 3-17.

Aerosols generated by flashlamp exposure to acetone vapor was found to scale exponentially with flashlamp fluence (Figure 3-21). During these experiments, the flashlamp drive voltage was not changed, and the varying fluence was achieved by turning off various lamps in a Shiva-size half-shell. This technique assures that the spectrum of the lamps did not vary over the range of values of the fluence output from 1.7 to 10 J/cm². As a rule, one should expect the volume of material ablating off a surface to be proportional to the lamp fluence incident on that surface. Since pulse duration was fixed for these tests, the Stephan-Boltzmann law should relate flashlamp plasma temperature to wall fluence. As such, a fourth-power relation between lamp energy and aerosol concentration should be expected. The lamp spectrum is not really black-body, and the lamp envelopes are doped with ceria to suppress UV, so it is hardly surprising that the simple relation is not quite observed; rather than a fourth-power law, the dependence is closer to the 3.75 power.




Figure 3-21. Aerosol particle concentration as a function of flashlamp fluence. Concentration increases exponentially with an exponent of 3.75.

3.7.4.6 Correlation between Aerosols and Slab Damage

Currently, there is no direct evidence linking slab damage to settled aerosol particles. However, aerosol particles placed directly onto glass samples have resulted in damage after a single flashlamp firing, see Section 3.7.2.2.

The general hypothesis is that a film of NVR averaging about 10–20 nm thick “remains on” amplifier surfaces after cleaning and spalls off when exposed to flashlamp light. The result is a cloud of ~30-nm black, mostly hydrocarbon particles. This population forms the feedstock for carbonaceous aerosol. While this cloud may be relatively hot, partially oxidized with the low O_2 concentration available in the predominantly N_2 atmosphere and even electrically charged, we are interested in how it behaves several seconds later when it has cooled, has reached a steady state, and is electrically neutral.

Somehow, shortly after spalling off, these small particles coagulate to form larger “smoke” particles that look like low-density aerogel foam. Photographs of some of this smoke can be found in Figure 3-19. Much of the smoke seems to consist of loose conglomerations of as many as 100,000 feedstock particles. Interspersed in this population are occasional higher-density blobs that appear to be sintered lumps of the same or similar hydrocarbon materials. Direct observation, as well as settling rate

measurements on the largest smoke “dust-bunnies,” establish their density as about 0.1 to 0.2 gm per cc or about 1/10th solid carbon density.

Let N_k be the number of smoke dust-bunnies per unit volume made up of k particles. Assume that all collisions are inelastic and that equi-partition of energy exists in the smoke cloud. Although the smoke is composed of low-density material, it seems reasonable to assume that the coagulated particles are roughly spherical, and consequently, several scaling relations emerge:

$$\text{diameter} = a_k \propto k^{1/3}$$

$$\text{area} = A_k^2 \propto k^{2/3}$$

$$\text{mass} = m_k \propto k$$

$$\text{velocity} = v_k \propto k^{-1/2}$$

The time rate of change of N_k depends on how many k -particle dust-bunnies are created and how many are lost per unit time. The letter p is a rate constant.

$$\frac{dN_k}{dt} = \frac{1}{2} \sum_{i=1}^k p_{i,k-i} N_i N_{k-i} - \sum_i p_{i,k} N_k N_i$$

The first summation denotes the number of k element smoke particles created per unit time by the collision of two smaller particles whose masses add to m_k . The second, a loss term, represents collisions in which members of the k bin collide with any other species and thus leave the k bin. There must also be a conservation relation since none of the feed-stock particles are created or destroyed following the initial generation event. If the total number of small particles is designated N_0 , then

$$N_0 = \sum_i i N_i$$

We are principally concerned with the largest smoke particles as these are simultaneously the most dangerous and the rarest. If we specialize to the largest particles, then k assumes the largest value (several hundred thousand perhaps), $N_k \sim 1$, and $N_i = 0$ for $i > k$. This lets the second summation stop at k . Also, approximate steady-state conditions are sought, even though formally this system can never reach steady state in any finite time, so we specialize to late times when the time derivative of N_k is small. One way for the summation on the right-hand side above to equal zero would be for each element to individually vanish. This leaves us with an approximate relation between the bin population densities:

$$N_{k-i} = N_k \frac{2p_{i,k}}{p_{i,k-i}} .$$

Heuristically, it is reasonable that the rates depend on both collision partners existing in the same location during the interval dt . The probability that this is true scales (by analogy with colliding beams) as:

$$p_{i,k} N_i N_k = \text{Const.} \frac{N_i v_i}{a_i^3} \frac{N_k v_k}{a_k^3}$$

If k is the largest value, then $k-i$ covers all integers from 1 to k . Canceling all common factors yields an approximate power law for any bin density

$$N_i = \frac{\text{Const.}}{a_i^{4.5}} \quad .$$

However questionable the derivation, this power law resembles what is actually observed. Figure 3-17 shows cumulative particle counter traces, $\int_a^\infty N(x) dx$, obtained in August of 1998 during a series of flashlamp shots. These data seem to be typical of such traces, and when replotted as a function of particle diameter, display a negative ~ 3.5 slope, as shown in Figure 3-20; therefore, the distribution function falls off as $\sim a^{-4.5}$. As so often happens, this problem is very similar to a famous derivation by Chandrasekhar found in “Stochastic Problems in Physics and Astronomy,” *Reviews of Modern Physics*, Vol. 15, Number 1, January 1943, pp. 59–63.

Chandrasekhar’s famous analysis builds on even earlier work by Smoluchowski on coagulation. Chandrasekhar begins by analyzing the case of a particle in a medium of infinite extent in which a number of similar Brownian particles are distributed uniformly at $t = 0$. He then calculates the rate at which these Brownian particles arrive at the sphere of radius $R = a$, surrounding the fixed particle. The general solution is written in terms of a normalized time:

$$\tau = 2/3 \frac{k_B T}{\eta} t \quad ,$$

where η is a frictional viscosity coefficient, k_B is the Boltzman constant, and T is the absolute temperature. The elegant closed solution to the family of rate equations becomes tractable due to a property of the diffusion coefficients (given that Brownian motion is the main effect), and the result is

$$N_i = N_0 \frac{(N_0 \tau)^{i-1}}{(1 + N_0 \tau)^{i+1}} \quad .$$

Here k can be large; e.g., 100,000, and N_0 is also large; e.g., $\sim 1,000,000$ particles per cubic foot. The relevant range of interest for these parameters seems to be $1 \ll N_0 \tau \ll k$. In this range,

$$N_i \approx \frac{1}{i\tau} = \frac{1}{Tt} \frac{\text{Const.}}{a_i^m}$$

where m is a power law (typically >3). We should expect an inverse cubic or greater scaling with particle size as well as an inverse relation with both time and temperature. At very early times there can only be population in the lowest bin and conversely at infinite time, all the particles must have coagulated into a single entity. Our particle counter and our damage data are both presented as “cumulative distributions;” i.e., the number of particles or sites “above a certain size.” This has the effect of reducing the logarithmic slope by one from the power of the density distribution. Thus, a negative fourth-power density-function scaling would appear as approximately an inverse cubic cumulative distribution.

The most consistent data set available so far seems to have emerged from Beamlet measurements taken just as that laser system was being decommissioned. Aerosol measurements done on 1/12/98 showed that just following a shot, the amplifier environment was at Class 200,000, and the distribution function, $N(a)$, could be fit by

$$N(a) = \frac{4072}{a^{4.24}}$$
 for a in micrometers. The next shot came about 200 minutes later, by which time a combination of clearing and sedimentation had reduced the amount of smoke to the point that the aerosol particle distribution could be fit by

$$N(a) = \frac{3751}{a^{4.846}}.$$

Given a smoke-particle-size distribution, we can begin to explore its relation to the observed damage. Presumably those smoke particles that are large enough to absorb flashlamp emission efficiently ($a > 5$ nm) and too large to lose their heat to the gas around them, try to come into black-body equilibrium with the lamps. To do so they would have to reach about 10,000 degrees, which turns them into strong UV radiators and often into plasmas. Surely, those smoke particles that happen to lie within their diameter of the laser glass surface at shot time can be expected to mar the glass surface.

After disassembly, Beamlet amplifier slabs were each scanned using a flat-bed optical scanner. Feit averaged these data (see Section 3.7.6) and arrived at a best-fit power law for the scanner’s data, assuming that the damage sites could be represented

by circles, after integration to extract the distribution $N_D(d) = \frac{15720}{d^{3.8594}}$. He also found a

fit assuming elliptical damage sites with an aspect ratio of about 4. For this fit, when plotted against half the long dimension, the multiplicative coefficient becomes 114,100, but the power law remains unchanged. Figure 3-14 displays a photograph of lamp-induced damage. A single lamp illumination produces “flash-burn” crazing of the surface near small particles placed on a BK7 glass plate. Aspect ratios between one and two seem typical of such burns. As Figure 3-15 shows, an approximately linear relation is observed between the size of the dust particle and the resulting damage, $d \sim 7.7 a$. No data has been presented concerning how such damage sites may grow when exposed to additional lamp illuminations, however, if this damage is at all like laser-induced damage, we can expect at least some growth. Smoke-particle-size density distributions should be one of several factors contributing to observed damage distributions. Given unblemished slabs going into the amplifier, particles either settle on or are merely near the glass surface when the flashlamps fire, and flashburns are generated. A procedure

to estimate the damage distribution given the smoke-particle distribution is outlined below:

- Number of particles of diameter 'a' per unit volume = $N(a)$
- They can reside in $\frac{\text{Amplifier volume}}{a^3}$ locations
- Assuming 'dangerous' sites must lie within 'a' of a glass surface, the number of risky locations is $\frac{\text{Glass area}}{a^2}$
- The fraction of the particle locations that are risky is F_r

$$F_r = \frac{\text{Glass area} \times a}{\text{Amplifier volume}} \sim 5.5 \times 10^{-6} \text{ for } a \text{ in } \mu\text{m}$$

Then, the number of damage sites of dimension 'd' counted after N shots should be greater than

$$N_D(d) > N_s F_r N(d/7.7 \times \text{growth factor})$$

We can apply this method to the Beamlet data set. The *Glass area* in the NIF power amplifier is about 280 sq ft., while the *Amplifier internal volume* is about 170 cu ft., which establishes a value for F_r . N_s , the number of shots, is 1500, therefore

$$\frac{15720 \rightarrow 114100}{d^{3.8594}} \geq 1500 \times 5.5 \times 10^{-6} \times \frac{3751}{\left(\frac{d}{7.7G}\right)^{3.846}} = \frac{79440}{d^{3.846}} \times G^{3.846}$$

Apparently the growth factor, G , must be near unity and the power law which emerges is very nearly the same, $-3.8594 \sim -3.846$.

Given a damage distribution function, we can compute the obscuration area per cm^2 of each slab surface in square microns,

$$A_o = \int_{25}^{1000} \pi r^2 N_D(r) dr = 3463 \mu\text{m}^2$$

and this implies an obscuration fraction per surface of 3.5×10^{-5} (or twice this value per amplifier slab). This number is heavily influenced by small scatterers at the limit of the scanner's detection capability, so the value above may well be low by a factor of two. In short, if carbonaceous aerosols are present at shot time, we can expect slab damage with a distribution related to the aerosol distribution. Such damage should continue to occur so long as the aerosol persists, hence sweeping out the "smoke" between flashlamp firings as Nova does, is the proper course of action. This mechanism having been thwarted, however, others emerge as we will see below.

As a practical matter, there is no clear-cut demonstration that damage can be turned off or on with the absence or presence of aerosols. Experiments to simulate an aerosol-free environment are explained in Section 3.7.10.

3.7.5 Slab Damage Analysis and Propagation Modeling

3.7.5.1 A Slab Damage Criteria

To establish refurbishment rates and define possible downstream damage, a method is needed to quantify damage, and a criteria is needed for the upper limit to the quantity of damage. To establish a useful lifetime, a criterion such as the following needs to be established:

Laser slabs shall be capable of >10,000 shots (at full laser intensity) without causing “end-of-life” damage to themselves or to other optics in the chain.

The quantity of damage can be established as a maximum damage spot size per surface or as a maximum scattering per surface.

Widmayer⁹ has used a sophisticated modeling code to predict the peak intensity on spatial filter input lenses based on small obscured spots at various locations on laser slabs. Results shown in Figure 3-22 indicate that defects as large a 1-mm result in beam modulation (ratio of peak to mean intensity) of 1.75, where this value is a factor of 2.7 below the laser damage threshold of the spatial filter lens.




Figure 3-22. Predicted relative intensity on the cavity spatial filter input lens due to small opaque defects of various sizes on the slabs preceding the spatial filter input lens. The presence of even a 1-mm size defect is safe and will not endanger the safety of the integrity of the cavity spatial filter input lens.

⁹ Clay Widmayer and Paul Renard, “Survey of damage threat to vacuum barriers by defects and inclusions in NIF optical components,” NIF-0008158, March 15, 1998.

3.7.5.2 *Damage Size Distribution*

The cumulative size distribution of the damage (or surface obscurations) on slab surfaces has been measured many times, and it is found to have a slope of -2.5 , which is very similar to the cumulative size distribution observed in Fed Std 209 and to the size distribution of the aerosols that appear after a flashlamp firing (see Figure 3-20).

Damage measured in AMPLAB experiments was observed to increase slowly with the number of shots over the range from 220 to 750 shots. However, the slabs in these tests also had an unusually high level of initial surface damage. Figure 3-23 shows the cumulative size distribution of laser slabs tested during the AMPLAB experiments. The damage concentrations observed on AMPLAB are typically $10 \text{ obscurations/ft}^2 \geq 100 \text{ }\mu\text{m}$.

Figure 3-24 shows the cumulative size distribution of laser slabs removed from Beamlet after 1500 shots. The damage concentrations observed on Beamlet are typically $40\text{--}50 \text{ obscurations/ft}^2 \geq 100 \text{ }\mu\text{m}$.

Figures 3-31 and 3-32 show similar curves for slabs removed from Nova when it was disassembled. The damage concentrations observed on Nova are typically $200\text{--}400 \text{ obscurations/ft}^2 \geq 100 \text{ }\mu\text{m}$ and are therefore nearly $10\times$ higher than damage observed on AMPLAB.





Figure 3-23. Cumulative size distribution of damage sites on slabs removed from AMPLAB. Concentration at $\geq 100 \mu\text{m}$ is approximately 10 obscurements / ft^2 .

In general, all Beamlet slabs had less damage than those removed from Nova and more damage than the slabs removed from AMPLAB. It is generally believed that the amount of damage is also a function of the number of flashlamp shots, and Nova had substantially more shots than either Beamlet or AMPLAB. However, controlled experiments to demonstrate the growth of damage as a function of the number of flashlamp shots will be done in FY2000.




Figure 3-24. Cumulative size distribution measured on slabs removed from Beamlet when it was disassembled. The concentration at $\geq 100 \mu\text{m}$ is approximately 40–50 obscurations/ ft^2 .

3.7.5.3 Beamlet Damage Has Been Measured with Two Techniques

Slab damage has historically been measured by experienced technicians using a high intensity light and a reference scale. Recently, this has been supplemented by the use of an optical flatbed scanner. When identical slabs were scanned and manually counted, it became apparent that the manual counting generally underestimated the smaller-sized damage sites and overestimated the few larger damage sites. Because person-to-person variation can also be eliminated by the use of a flatbed scanner, all future slab damage counting will be done with some form of automated counting system.

3.7.5.4 Beamlet damage results compared to AMPLAB

Although Beamlet and AMPLAB amplifiers were similar in construction, several second-order effects may help explain the higher damage observed on Beamlet. These effects include:

- Epoxy blastshield seals on Beamlet were exposed to flashlamp light and were severely burned, presumably resulting in the generation of considerable aerosol.
- Beamlet slabs received 1500 shots versus only 750 on AMPLAB.
- Initial Beamlet blastshields were sol-gel coated while mounted in their metal frames and the sol gel failed to adhere to the metallic surfaces and resulted in significant infant mortality to the laser slabs.

3.7.6 Cumulative Average Damage Distribution for Beamlet

Feit¹⁰ has studied the cumulative size distribution and made adjustments for the shape of the damage sites, which acknowledges that the downstream damage to optics is dependent on the maximum size of damage sites rather than simply their average diameter. The slope of the size distribution analyzed by Feit is -2.8 , which is larger than previously reported (see Figure 3-25).

The cumulative size distribution extrapolated from the Beamlet slab damage data can be interpreted as a failure rate for NIF. Figure 3-26 indicates that depending on the shape of the damage (circular or elliptical with aspect ratios of 2:1 to 4:1), there could be slab damage of 2-mm size or greater on each surface of 10 to 35% of the NIF slabs after only 1500 shots (two years).




Figure 3-25. Cumulative size distribution of slab damage measured on Beamlet after 1500 flashlamp shots. The curves include adjustment for the shape of slab damage for an aspect ratio of 4:1.

¹⁰ M. D. Feit, "Analysis of Beamlet amplifier slab damage distribution data," NIF-0021664, May 3, 1999.




Figure 3-26. Cumulative probability of damage based on Beamlet slab damage data after 1500 shots. Easily 20–30% of the NIF slabs could sustain critical size flaws after only 1500 shots based on the precise shape of the damage and the specific maximum allowable damage size.

3.7.7 Logic and Estimated Cost of Slab Refurbishment

3.7.7.1 Considerations for Slab Refurbishment

Spaeth has considered the consequences and costs associated with amplifier slab refurbishment. The first consideration is whether simple slab cleaning could significantly reduce surface obscurations and therefore negate the need for more costly regrinding and polishing. Although it has been observed that as much as half the obscuration on slabs is due to opaque but removable contaminants, it has also been observed that when a damaged slab is recleaned, it damages again very quickly after only a few subsequent flashlamp shots. One AMPLAB slab was removed and recleaned, and when reinstalled, its surface damage (at the $\geq 500 \mu\text{m}$ size) increased from $0.3/\text{ft}^2$ to $4/\text{ft}^2$ in only 37 shots. Therefore, until further information is obtained, it is not recommended that slabs be removed, cleaned, and reinstalled, as this is likely to result in rapid and significant surface damage.

If slabs are removed for refinishing, the existing damage sites must be removed by grinding and polishing to a depth that is at least the depth of the damage and of the

underlying cracking. Slab damage depth is observed to be about 10% of its diameter but it is also generally known that subsurface cracking extends to a depth of several times the size of the abrasive used to remove the damage. With damage diameter up to 1–2 mm, the initial subsurface damage can be expected to extend at least 400–600 μm , which may be removed from both sides of the slab. A review of slab dimensions (40- to 41.5-mm thickness specification) and an observation that the current LRU can only accommodate ≥ 37 mm slabs, has led us to conclude that any given slab can be repolished at most two times.

3.7.7.2 Costs Considerations for Slab Refurbishment

Costs parameters include the initial slab cost based on the new continuous pouring process plus the cost of edge cladding and polishing. The cost savings associated with the new continuous pouring process has reduced the glass blank cost by a factor of 2 \times while other finishing costs have not changed significantly and are similar to prices paid for Nova slabs (adjusted for inflation). Current plans call for the procurement of about 10% construction slabs.

The integrated cost of slab refurbishment is therefore highly dependent on the rate that slabs exceed their maximum damage criteria (2- to 3-mm maximum size damage) and the quantity of spare slabs available that can be substituted for damaged slabs. For example a 9% failure rate would initially be resolved by replacement with polished slabs purchased as part of the construction spares. Subsequently, the damaged slabs would be repolished at moderate cost until they become too thin for reinstallation. At this point, new glass slabs will need to be purchased. However, if the continuous pouring factories have been shut down, the only apparent solution is to purchase conventionally fabricated glass at a significantly higher cost. The 9% failure rate previously mentioned would asymptotically result in the procurement of $9\% \times 3000$ or 270 slabs/yr for a sustained annual cost of \$10–\$12M.

3.7.8 Contaminant Mitigation, Design Strategies, and Operational Strategies

3.7.8.1 Components of a Candidate Risk Mitigation Strategy

A solution to the laser amplifier slab damage problem is clearly dependent on the acquisition of better damage data through the use of a small-scale test facility, the testing of full-size slabs as in an AMPLAB II experiment, and the collection of further damage data from Nova and Beamlet slabs. Furthermore, there are several possible designs and operational strategies: improved cleaning processes to remove residual organic films from structural surfaces, elimination of elastomers within the amplifier cavities, and development of gas flow scenarios that will flush the aerosols from the slab cavity. And finally, development of new modeling and analysis tools to understand the initial slab damage mechanism.

3.7.8.2 Possible Responses to Reduce Damage from Aerosols

Improved cleaning of enclosure walls will be provided by industrial experts who utilize a high-pressure spray water washing system incorporating a demonstrated surfactant. The detergents from the Brulin Corp when applied hot and thoroughly rinsed with high-pressure DI water at 3500 psi or washed in a high-energy-density ultrasonic cleaner have shown that they can consistently produce surfaces with organic levels of $\leq 0.1 \mu\text{g}/\text{cm}^2$, which is equivalent to only a few monolayers of organic material.

Flashlamp cleanup of aerosols was first observed during the AMPLAB experiments and provides a mechanism for reducing organic films by allowing them to be pyrolyzed after the amplifier is completely assembled, but before the laser slabs are installed. Commissioning activities will include ≈ 300 monitored firings to watch the progressive reduction of aerosols before the slab cassette is installed.

The elimination (or shielding) of elastomers from any region where they would see high-fluence flashlamp light is planned as well as the use of elastomers that are inherently low in aerosol production. Extensive small-scale experiments have shown that elastomers should be first selected by exposure to full-fluence ($10 \text{ J}/\text{cm}^2$) flashlamp light and then evaluated for aerosol production at various reduced flashlamp fluence levels achieved by turning off various fractions of the flashlamps. The aerosol production should then be scaled to produce an aerosol of only a few percent of the expected aerosol from the entire amplifier. Current plans are that after the 300-shot burn-in, an amplifier should be producing a peak aerosol of between Class 1000 and Class 10,000 and that it should remain in the amplifier for less than 10 minutes.

Sections 3.7.8.4 through 3.7.8.5 describe a plan to divert a portion of the flashlamp cooling gas and allow it to flow through the slab cavity to flush out aerosols that form immediately after a flashlamp firing. This gas flushing will only be activated for a short duration (30–60 minutes) shortly after an amplifier is fired to dilute the stagnant gas within slab cavity.

3.7.8.3 Amplifier Elastomer Status

A design conflict exists within the slab amplifiers concerning the use of elastomers to allow gas volume pressurization. To maintain a positive pressure within the slab cavity (and within other interconnected cavities adjacent to the Main and Power amplifiers) a seal must exist at every assemble joint. Many of the structural joints will be welded, and some will be fitted with flat silicon gaskets. Within the amplifier slab cavity, the walls are either silver reflectors or open to the next slab cavity or blastshield glass. The blastshield glass must be gas tight, as high-velocity cooling gas needs to flow through the flashlamp cavity, and this requires the flashlamp cavity to operate at a higher average pressure compared to the slab cavity. An all-metal, glass-to-metal seal might be possible to design, but a considerably less expensive solution is the use of a cast elastomer gasket. Since most elastomers are easily pyrolyzed by the high-fluence flashlamps, it is imperative that the gasket be protected from direct flashlamp and laser light illumination.

On Nova, there are no seals around the blastshields, as both the slab and flashlamp cavities are connected to a common gas supply plenum. Although the Nova gas supply system was designed to have a high-flow-rate flashlamp cooling mode as well as a low-flow-rate beamline pressurization mode, it was never operated in the high-flow-rate mode.

3.7.8.4 The Historic and Current Views of Gas Flow

The historic view of pressurized laser cavity systems has been to enclose all precision-cleaned components and laser optics within a slightly pressurized beamline cavity. The clean gas over-pressure allows gas to flow outward and thereby prevents the relatively dirty environmental gas of the Laser Bay from flowing back into the laser cavity. The quiescent aerosol in the laser cavity is generally below Class 1, whereas the Laser Bay is nominally Class 100,000.

Since the first measurement of flashlamp-generated aerosols within an amplifier flashlamp cavity on Beamlet, the understanding of the purpose of a protective gas has changed. The generation of aerosols thousands to hundreds of thousands of times the ambient aerosol level (sometimes approaching Class 1,000,000) is contrary to maintaining a clean environment around the optical components. Current NIF cleanliness requirements state that optical surfaces will be cleaned to Level 50, and they should be removed for refurbishment when the particles or surface damage reaches 2.5×10^{-4} , or when any 2-mm damage is observed on the surface. An optical surface cleaned to Level 50 will require only 100,000 Class-hours of exposure to reach Level 125. Correspondingly, to extend slab cleanliness lifetime (time between refurbishment) to 10 years (7500 shots or 87,600 hours) will require lowering the peak aerosol concentration to below Class 100 and the average aerosol below Class 1. Flowing gas through the slab cavity immediately after a flashlamp firing lowers the peak aerosol concentration and also lowers the integrated exposure time.

3.7.8.5 The Current View of Slab Cavity Gas Flow

Two geometric options exist to achieve rapid slab cavity gas flow: horizontal flow from one end of the amplifier to the other while wandering around the slabs mounted at Brewster's angle; and vertically with gas flowing parallel with the slab surfaces. The horizontal flow configuration would be easier to implement in the existing laser chain configuration, but as shown in Figure 3-27, the wandering horizontal flow transports airborne particles on a trajectory directly past slab surfaces and with up to 11 slabs in serial configuration, this would significantly increase the probability of the aerosol depositing onto slab surfaces.

The vertical flow configuration flushes the cavity with a significantly lower probability of collisions between the aerosol particles and the slab surfaces. After evaluating many alternatives, a configuration has been chosen that diverts 10% of the flashlamp cooling gas through a port in the top of the amplifier, directs the flow around the upper silver-coated reflector and past the four slabs, and exits the gas around the

lower reflector and out through a cavity beneath the amplifier. The vertical flow will correspond to 30% of the amplifier volume per minute, which will lower the aerosol concentration to 5% of its initial value in 10 minutes.




Figure 3-27. Streamlines of airflow around laser slabs configured for horizontal flow. This configuration unfortunately results in a relatively high probability of collisions of the aerosol particles with the slab surfaces, which is nearly completely avoided in the vertical flow configuration that was eventually chosen for implementation.

3.7.9 Retrospective on Nova Amplifier Slabs

3.7.9.1 Selection of Amplifiers for Examination

Prior to decommissioning Nova in 1999, the input and output nitrogen flow rates were measured on all of the slab amplifiers. In addition to the flow-rate measurements, accurate records of flashlamp firings for each slab amplifier were researched. The amplifiers were then categorized into four groups; high flow rate and high shots (HFHS), high flow rate and low shot (HFLS), low flow rate and high shots (LFHS) and low flow rate and low shots (LFLS). All of the Nova slab amplifiers (46-cm, 31.5-cm, 20.8-cm, 15-cm, and 9.4-cm) were categorized in this fashion to correlate with slab damage. It had long been believed that slab damage increased linearly with flashlamp firings, and although a positive flow of nitrogen was believed to exclude contamination, no data had ever been collected to correlate slab damage with the amount of flow.

3.7.9.2 Slab Damage Data Acquisition

The damage data was obtained using a 800-dpi flatbed scanner and data acquisition software developed by Doug Ravizza. The combination of the scanner's short depth of field and the relative opacity of the laser glass permits the scanning of one slab surface at a time. The scanner does not distinguish between damage and an obscuration due to

a surface dust particle. Great care was taken to insure that the illumination was identical for all slabs. The image file data was reduced by thresholding the grayscale file and then by determining the enclosed pixels from which an equivalent circular diameter is calculated. In other words, for an obscuration to be counted, it must be above a certain grayscale level (predetermined and identical for all slabs) and it must have an area of ≥ 3 pixels. All other data are ignored. The smallest obscuration diameter that can be seen using this procedure is $62\text{ }\mu\text{m}$.

To verify the validity of this technique, a study was done comparing the flatbed scanner results with the results of four Nova technicians who assessed and categorized damage “by eye”. The scanner technique was in agreement with the “by-eye” technique for large obscurations ($\approx 1\text{ mm}$) and yielded many more obscurations per unit area for smaller obscurations ($\approx 100\text{ }\mu\text{m}$). A final calibration of the flatbed scanner technique used a calibrated mask imprinted with well-defined shapes and sizes around $200\text{ }\mu\text{m}$. The flatbed-scanner technique acquired and sized all features correctly.

3.7.9.3 Obscurations Maps

A typical obscuration site map for a Nova disk is shown in Figure 3-28. Each obscuration location is marked by a spot, where the spot size (not to scale) is indicative of the obscuration size. A histogram of these obscurations vs. obscuration diameter is generated for each slab surface. The obscurations are normalized to obscurations/ ft^2 and integrated. The results are then presented in a form shown in Figure 3-29, where the cumulative obscurations/ $\text{ft}^2 \geq$ diameter of the obscuration are plotted against the obscuration diameter on a log-log scale.




Figure 3-28. Obscuration site map for a typical Nova 31.5-cm disk showing the location of the obscurations and their relative sizes.




Figure 3-29. Cumulative obscuration density plot of Nova disk SP4-416 showing cumulative obscuration densities vs. diameter.

3.7.9.4 Nova 31.5-cm Amplifier Slab Damage Data

A schematic of the Nova 31.5-cm beamline is shown in Figure 3-30. Each amplifier contains two slabs, and five of these amplifiers constitute a beamline and flow system. Two such beamlines were analyzed: beamline 1 (LFLS) and beamline 10 (HFLS). The cumulative obscuration data from both beamlines are shown in Figure 3-31 and 3-32 respectively. Although the damage data can vary by as much as 10 \times , the beamline-to-beamline variations are small and the shapes of the curves are very similar. For these reasons, one can use the value of a given curve at a diameter of 100 μm as representative of the entire curve to facilitate further data reduction and analysis.




Figure 3-30. Top view of a Nova 31.5-cm amplifier beamline.




Figure 3-31. Cumulative obscuration density for Nova amplifiers on Beamline 1.




Figure 3-32. Cumulative obscuration density for Nova amplifiers on Beamline 10.

A plot of the obscurations/ $\text{ft}^2 \geq 100 \mu\text{m}$ vs. number of flashlamp shots is shown in Figure 3-33. One immediately notices that there appears to be no correlation between the number of shots and obscuration density. Similarly, if one plots obscurations/ $\text{ft}^2 \geq 100 \mu\text{m}$ vs. flow rate (Figure 3-34), one sees no correlation between flow rate and obscurations. These results lead one to conclude that damage on Nova slabs is based on some initial condition and does not appear to increase after 200 shots. (There are no Nova damage data at fewer than 200 shots, and there are no pristine Nova slabs to which we can compare these results.)




Figure 3-33. Obscuration density plotted vs. number of shots for two of the Nova 31.5-cm beamlines.




Figure 3-34. Obscuration density vs. gas flow for two of the Nova 31.5-cm beamlines.

3.7.9.5 Mechanical Damage Observed on Nova Slabs

Upon close examination of Figure 3-28, one notices that there are more obscurations near the edge of the slab, which is nominally covered by a mask, compared to the number of observations in the center of the slab. A plot of obscuration density vs. normalized radius for the Nova 31.5-cm beamline 1 (Figure 3-35) shows a significant ($\sim 2\times$) increase in obscuration density near the mask edge. An examination of one of the 31.5-cm masks (Figure 3-36) shows sharp and burred edges on both the step edge and the mask edge as well as a direct mask contact onto the laser glass due to an inherent overlap in the mask design. During a flashlamp pulse, an acoustic shock is generated that shakes both the mask and the slab. This shaking (or vibration) can lead to the mask contacting the slab surface, chipping the glass, and increasing the obscuration density seen near the mask edge. Mechanical contact with the pumped gain medium can lead to damage in another way as well. Significant fluorescent flux, over 0.3 J/cm^3 , exists inside the pumped laser slab. Metal contact points on the slab periphery are within the fractional wavelength limit required for efficient evanescent wave coupling. This mechanism would tend to locate damage sites at the mask edge on the lower surface of the slab. Molten metal splashed by this mechanism could account for lower surface damage site concentration falling off as it moves away from the mask contact points. The location and concentration of Nova slab damage sites argues against faulty cleaning methods, large aerosol levels, or laser beam-induced effects and for mechanical or evanescent wave coupling at the lower slab edges. These problems will be avoided in the NIF by contacting the vertically mounted slabs well away from the illuminated region on the edge cladding only.




Figure 3-35. Obscuration density vs. normalized radius for Nova 31.5-cm beamline 1.



Figure 3-36. Photograph and schematic of Nova 31.5-cm slab holder. Notice the regions where the step edge and mask edge overlap.

3.7.9.6 Removable Obscurations

Of the forty 31.5-cm surfaces (two beamlines, ten slabs/beamline, two surfaces/slab) that were examined, four surfaces were chosen to be examined before and after precision cleaning. The cumulative obscuration density plot for one such surface is shown in Figure 3-37. A $\geq 2\times$ reduction in obscuration density was observed on all four slab surfaces after cleaning. This leads to the questions, “what were these obscurations and why did they not cause permanent slab damage?” Contaminant samples were taken from several random locations for SEM examination. The SEM results indicated that the majority of the removable obscurations were SiO_x . From this observation, it appears that while SiO_x particles may scatter light out of the laser aperture, they do not seem to cause damage to the amplifier slabs.




Figure 3-37. Obscuration density on a Nova 31.5-cm slab, before and after precision cleaning.

3.7.10 The Small-Scale Flashlamp Damage Testing

3.7.10.1 Goals for the Small-Scale Test Facility

The goal of the small-scale test facility is to investigate and determine the sources and causes of laser glass damage due to flashlamp light. A partial list of possible sources of contaminant include: airborne and settled aerosol particles, NVRs on surfaces, elastomer seals, and cleaning residue. A partial list of possible causes of damage include: differential cooling between laser glass and contaminant, evanescent wave coupling between contaminant and laser glass, and surface defects in laser glass.

Experiments were performed by placing a variety of airborne particles onto laser glass and exposing them to flashlamp light. This type of experiment nearly always results in a slab damage spot even after a single flashlamp shot. Despite our confidence with this direct aerosol damage demonstration, a direct observation of the generation of an aerosol by the pyrolysis of organic films on structural surfaces, the transport of the aerosol to the glass surface, and the subsequent damage to the glass has not been made. That is, we postulate that a flashlamp exposure test resulting in the formation of an aerosol should always lead to slab damage, and a corresponding experiment in which no aerosol is observed should result in negligible slab damage. As simple as this experiment sounds, it has proven impossible to construct a sufficiently clean vessel that does not generate an aerosol when exposed to flashlamp light and correspondingly does not lead to slab damage on the surface of a glass sample within the vessel.

Results from Nova (described in more detail in Section 3.7.9) indicate that there appears to be an “infant mortality” phenomenon in laser slab damage. After some

number of shots, there is no correlation between the laser slab damage and number of shots. The small-scale test facility will allow us to examine the laser slab after any fixed numbers of shots to determine the correlation between shot number and damage.

The small-scale test facility also permits the investigation of “designer” aerosols and contaminants (with a specific constituency and density). Preliminary results suggest that some constituents (silica) do not result in slab damage. The small-scale test facility will allow the experimenter to examine the types of cleaning and surface preparation methods that may mitigate damage.

All of the above experiments are much more easily performed in a well-controlled, small-scale environment than in a full NIF-scale environment. Recent studies have shown that the results from these small-scale experiments are, in fact, scalable to Nova- and NIF-size laser slabs.

3.7.10.2 Experimental Goals of the Small-Scale Flashlamp Damage Test Facility

The overall goals of the small-scale flashlamp damage testing facility are to demonstrate:

1. A direct correlation between aerosol production rate and damage rate.
2. A measurement of the rate of slab damage vs. the number of flashlamp shots and a correlation with initial cleanliness and aerosol generation rate.
3. The ability to distinguish between intrinsic and extrinsic slab damage (i.e., is the damage caused by initial surface defects or entirely related to aerosol deposited initiation sites).
4. Determination of the size of initial surface defects that map into final surface damage. This would establish a minimum-size defect to be measured during subsequent slab cleaning and slab installation activities.

3.7.10.3 Features of the Small-Scale Test Chamber

A flashlamp test chamber was constructed to simulate the environment inside of a laser slab amplifier. The chamber was designed to fit inside the Shiva 15-cm laser amplifier flashlamp half-shell used for materials testing. This half shell is mounted on top of a base that is 2.8 inches deep and 30×12 inches in length and width. The test chamber is the largest size that would fit inside the flashlamp half-shell, and its outside dimensions are $14 \times 11 \times 2.7$ inches (see Figure 3-38). The test chamber is constructed of sheet stainless steel that was bent and welded at the corners to form a nearly seamless enclosure. The top of the box has a flat sealing surface about 0.3 inch wide. Various flat gaskets were tried with little success in finding one that could withstand the flashlamp fluence without introducing large numbers of airborne particles into the test chamber. The seal that produced the fewest airborne particles, with an acceptable amount of leakage, was made from expanded teflon foam that was wrapped with aluminum foil to protect the teflon from flashlamp light. Flow measurements conducted on the chamber indicate that the leakage rate is about 10% of the inlet flow rate. The top cover of the test chamber is a 6-mm-thick plate of blastshield glass. The glass plate is pressed against the seal with a steel collar that is clamped with four bolts.




Figure 3-38. The one-square-foot box used for flashlamp exposure testing. The test box is mounted beneath a Shiva flashlamp half-shell.

The chamber is routinely cleaned using 10% isopropyl alcohol (IPA) and 90% water supplied in a premoistened clean-room wiper. We have also used methylene chloride and 100% IPA wipers for cleaning the chamber. In general, initial cleanliness as indicated by the airborne particle counter is nearly identical after any of these solvent cleaning methods. We observed a steady decrease in the initial airborne particle count with each additional shot for the first 10 to 30 shots. After this initial particle production, the box stabilizes at a flashlamp aerosol generation rate that is maintained for the next 100 shots. After cleaning the test chamber, we usually expose the empty chamber to 60 to 100 flashlamp shots to assure a relatively low background aerosol. The magnitude of the background is usually in the range of 10 to 100 particles $\geq 0.5 \mu\text{m}$ per shot. The gas flow rate through the test chamber and particle counter is $1 \text{ ft}^3/\text{minute}$, which means that the total particle production is from 10 to 100 particles emanating from the internal surface area of the test chamber. The chamber walls represent approximately 3 ft^2 , so the typical background generation rate is 3 to 33 particles/ $\text{ft}^2 \geq 0.5 \mu\text{m}$. When extrapolated to the surface area of a four slab amplifier, the particle generation rate per shot is expected to be 170 to 1,850 particles $\geq 0.5 \mu\text{m}$. This latter value corresponds to the generation of an aerosol density of only 77 particles/ ft^3 immediately after a flashlamp firing.

3.7.10.4 The Cylindrical Test Chamber with Metal Gaskets

To reduce the background particle count in the test chamber, a second chamber was designed that eliminates elastomer seals by using all metal seals commonly used by high-vacuum equipment. The test chamber, 8 inches in diameter by 2.5 inches high, is shown in Figure 3-39. Initially we planned to clean the chamber with IPA, seal it using metal gaskets, and bake the chamber to eliminate adsorbed organic residues on the chamber walls. Our initial vacuum window flange had a quartz window. During initial flashlamp exposures, we noted a change in the appearance of the quartz-to-metal interface and later learned that quartz windows are sealed using solder that melts during flashlamp exposure. The second window was made of pyrex with a kovar-graded glass-to-metal seal. This pyrex window has sustained more than 1000 shots with no noticeable degradation of the glass-to-metal seal.



Figure 3-39. The “bakeable” flashlamp exposure box used for flashlamp exposure testing. The test box is mounted beneath a Shiva flashlamp half-shell.

Initially the metal chamber parts were baked in an oven at 200°C with argon gas for 24 hours. When initially installed in the flashlamp half-shell, the test chamber was fitted with a nitrogen purge line and wrapped with heat tape to again bake out the chamber. The gas leaving the chamber was sampled with SPME (Solid Phase Micro Extraction) chromatographic probes. The SPME probe has an absorbent coating on a glass fiber that is exposed to the exiting gas. After exposure, the SPME probe is taken to a GC/MS chromatograph where the absorbed residue is thermally desorbed for separation (GC) and identification (MS). We found that during the first hour of chamber bake-out, organic materials were desorbed from the chamber walls. In the next two hours, no

further organic outgassing occurred. Analysis of the SPME probe revealed that the most common organic species was dibutyl phthalate (DOP) which is a common plasticizer. It appears that the chamber was an effective absorber of plasticizer that was naturally in the environment despite efforts to clean the chamber and protect it during transport and handling.

After the chamber was baked-out under nitrogen purge, we exposed the empty chamber to flashlamp light. The particle production by the baked chamber was essentially the same as the solvent wiped chamber. Generally, the initial particle generation rate is about 10,000 particles $\geq 0.5 \mu\text{m}$ diameter dropping to about 100, and occasionally to 10, particles per shot over a 20-shot sequence. The level of 100 to 10 particles $\geq 0.5 \mu\text{m}$ produced per shot seems to be the lowest that can be achieved with the current test chambers.

3.7.10.5 Laser Glass Sample Characterization

Prior to 1997, all previous documentation of laser glass damage was performed manually. That is, a high intensity light was used to illuminate the glass slab surface, which allowed visual counting of the damage spots. Both the size and the location of the damage spots were recorded on paper sheets referred to as damage maps. On a typical piece of Nova laser glass, approximately 100 damage spots / ft^2 could be located with a diameter $\geq 100 \mu\text{m}$. Since we are using laser glass samples of only one square inch, we realized that finding $\geq 100 \mu\text{m}$ damage spots will be unlikely. Therefore, we are developing a microscopic characterization method that will rapidly measure damage as small as $25 \mu\text{m}$ diameter. The characterization methods we tested to date have included direct photographic imaging, high-resolution CCD photography, and Schlieren illumination performed at low magnification and performed on a scanning imaging system to allow examination of a large fraction of the surface area of a glass sample.

The surface damage characterization method tested was to photograph the laser glass using grazing incidence illumination supplied by two fiber-optic bundles. A 35-mm camera with a 60-mm macro lens was positioned above the sample such that the sample nearly fully filled the photographic field. The film was developed and the images digitized at 3072×2048 pixels and stored on CD. The resolution using this approach was about $12 \mu\text{m}$ per pixel. The digitized image was read using Adobe Photoshop. Opaque objects on the glass sample surface were counted using the software Image Pro Plus. This photographic approach produced a permanent record but required several days to process the film and scan to a CD. These issues made the direct photographic approach very slow.

The second photographic method used a high-resolution CCD camera to photograph glass damage using the same grazing incidence illumination. The imaging system chosen was a Peltier-cooled Photometrics camera with 1640×1280 pixels. Further, the laser glass was mounted to allow one quadrant of the sample to be photographed as a single image, which resulted in a resolution of $10 \mu\text{m}$ per pixel. The cooled camera has very low background noise, allowing faint particles and damage sites to be recognized.

Analysis of the image for particles was done using either Matrox Inspector or Image Pro Plus. Unfortunately, we have found that the uniformity of illumination is poor, which makes it difficult to consistently adjust contrast and establish a threshold for particle counting. What is still needed is a very uniform illumination technique that enhances the contrast of the damage on the laser glass surface.

3.7.10.6 The Basic Laser Glass Flashlamp Damage Test Plan

Using the test chambers, we have planned experiments to investigate the mechanism of laser glass damage. The goal is to determine the effect of flashlamp fluence, glass cleaning methods, and the presence or absence of aerosol in the test chamber on the damage of laser glass. The fluence of flashlamp exposure is controlled by varying the number of flashlamps fired. The flashlamp half-shell contains six pairs of lamps wired in series, and it is possible to fire any combination of lamp pairs. The total fluence can thus be controlled from about 17% of full fluence to 100% (10 J/cm^2) in six steps.

The basic test plan involves initial characterization of the test chamber and the laser glass sample. The chamber is chemically cleaned, assembled, and exposed to flashlamp light until a stable background aerosol is achieved. The laser glass samples are chemically cleaned and characterized under a microscope before testing. The glass samples are then mounted into a holder and loaded into a test chamber (see Figure 3-40). The chamber is then exposed to flashlamp light for 60 shots. If a flow of $1 \text{ ft}^3/\text{min}$ is used, then the particle count is monitored continuously. If the flow is lower or turned off completely, then the particle count is monitored for only the first and last three shots. The laser glass is then removed from the chamber and characterized for damage.




Figure 3-40. One-inch-square by approximately 1/8-inch-thick phosphate laser glass test samples arranged to allow flashlamp exposure to both horizontally and vertically oriented samples. The pieces are cut from a larger slab of Nova amplifier glass and then ground flat. They are finally polished using cerium oxide abrasive for three days on each side.

3.7.10.7 Typical Experimental Results

Initial experiments were to establish the reproducibility of glass damage caused by flashlamp exposure. Glass slabs were cleaned using 10% IPA wipes and mounted in the test fixture as shown in Figure 3-40. All handling was done using Class 100 protocol inside a down flow clean room. The sealed chamber was transferred to the flashlamp facility and exposed to 60 shots with 1 ft³/min of nitrogen flow. All vertical slabs showed a damage region where the horizontal slab touched it. The slabs were stacked so that the horizontal slab supported the vertical slab. We did not include the damaged region on the vertical slab in our damage analysis. A region of interest was selected that excluded the bottom 4 mm that showed contact damage.

Precharacterization of the slabs showed that there was a range of initial cleanliness. However, the post-characterization showed consistently higher particle counts. In general, we found that the cumulative damage count for the glass slabs increased by a factor of 2 to 5 over the entire size range analyzed as shown in Figure 3-41.




Figure 3-41. Cumulative size distribution of the surface of several 1-inch-square glass samples exposed to 60 flashlamp shots.

Although we expected all damage counts in the reproducibility experiment to be identical, we found that the experiments fell into two distinct categories. In one group, the aerosol particle count from the chamber was about 1000 particles > 0.5 μm per cubic foot. The other group of experiments resulted in 100 particles > 0.5 μm per cubic foot. Inspection of laser glass from the two groups showed that there was more surface damage on the glass exposed in the higher aerosol particle environment.

A conclusive link between aerosol concentration and glass damage has not yet been proven. At this point we have established the methods for conducting and analyzing

the experiments. To verify that aerosol has a role in glass damage will require additional experiments with better environmental control. Modifications to the flashlamp exposure system have been made to provide characterization of the laser glass without handling or removal from the chamber. This should provide the added level of control needed to isolate the effects of initial cleanliness and aerosol concentration on glass damage.

4.0 NIF FINAL OPTICS ASSEMBLY AND 3ω OPTICS 4-1

Introduction..... 4-1

4.1 Frequency Conversion 4-3

References 4-7

4.1.1 Frequency Conversion Crystals; Growth and Finishing 4-7

4.1.2 Finishing Development 4-9

4.1.3 Crystal Fabrication Alignment 4-12

References 4-14

4.2 Lifetime of Optics in Operation..... 4-14

4.2.1 Reduction of Fused Silica Damage Initiation..... 4-16

4.2.2 Understanding and Reduction of Fused Silica Damage Growth 4-18

4.2.3 Mitigating Damage Initiation and Growth 4-21

4.2.4 Identification and Elimination of Mechanisms Leading to UV Damage of DKDP 4-23

References 4-30

4.3 Optics and Coatings 4-31

4.3.1 Improved Anti-Reflection Coatings for the Conversion Crystals 4-31

4.3.2 KDP Anti-Reflection Coating Development..... 4-31

4.3.3 Degradation of the Optical Performance of Sol-Coated KDP..... 4-35

4.3.4 Formation of Etch Pits..... 4-36

Engineering Model 4-38

Etch Pit Mitigation..... 4-39

4.3.5 Barrier Coatings..... 4-40

4.3.6 Coating Conclusions..... 4-44

4.4 Summary—NIF Final Optics Assembly and 3ω Optics 4-44

References 4-45

4.0 NIF Final Optics Assembly and 3ω Optics

Lloyd Hackel, Mark Kozlowski, Alan Burnham, Pam Whitman, Paul Wegner

Introduction

In the NIF design, the infrared light from the solid-state laser is converted to UV wavelength and spatially conditioned in the Final Optics Assemblies (FOAs). There are 48 FOAs attached to the target chamber, each configured to handle four beamlines in a square arrangement called a quad. Within each beamline is the 3ω optics assembly that includes a vacuum window, a potassium dihydrogen phosphate (KDP) frequency-doubling crystal, a deuterated potassium dihydrogen phosphate (KD*P) frequency-tripling crystal, a focusing lens, one or several diffractive optics plates and one or more glass plates, used as a shield for target debris. Figure 4-1 shows a picture of a prototype FOA, and Figure 4-2 shows a schematic layout of the FOA optics configuration of a single beamline.

The vacuum window serves as the barrier between the atmospheric pressure within the beam transport system and the vacuum of the target chamber. The pressure within the FOA can be maintained at vacuum or up to 200 Torr of a separate background gas. The doubling and tripling crystals convert the 1053-nm light to 351 nm. The diffractive optic plates (DOPs) are large optical gratings that perform three specific functions; the color separation grating (CSG) differentially diffracts the three wavelengths so that only the 351-nm light propagates to the target center. The beam sampling grating (BSG) diffracts a small portion of the UV light to an energy and power diagnostic. The kinoform phase plate (KPP) applies a spatially patterned phase modulation to the beam that, when combined with the frequency modulation of the input beam, smears out the speckle contrast of the output beam on target. Finally, the debris shield protects the optics from radiation, vapor, and debris ejected from the target. The greatest challenge to the FOA is to attain high-frequency conversion efficiency while maintaining low optical damage.



Figure 4-1. Prototype Final Optics Assembly undergoing functional tests. The target chamber vacuum isolation valve attaches at top, and the 1ω laser input enters the bottom of the assembly. Only one of four integrated optics modules is attached.

Figure 4-2 shows the optical schematic of the FOA.




Figure 4-2. Optical schematic of the FOA.

4.1 Frequency Conversion

The design of the NIF frequency converter has remained unchanged since the Title I review. Two crystals- an 11-mm thick Type I KDP doubler and a 9-mm thick Type II KD*P tripler- are arranged in quadrature such that the 1ω electric field is oriented along the ordinary axis of the doubler and along the extraordinary axis of the tripler. The 2:1 ratio of 2ω to 1ω energy at the input to the tripler that is critical for efficient mixing is achieved by angle-tuning the doubler $220 \pm 20 \mu\text{rad}$ from exact phasematching. The tripler is then aligned to its phasematching condition with a precision of $\pm 20 \mu\text{rad}$ to maximize the conversion efficiency. Because the tuning axes of the two crystals are mutually orthogonal, they can be aligned independently. Figure 4-3 shows the configuration of the frequency converter.




Figure 4-3. Basic configuration of frequency converters showing potassium dihydrogen phosphate Type I, second-harmonic generator (SHG) and deuterated potassium dihydrogen phosphate Type II, third-harmonic generator (THG).

Experiments performed at large aperture on the Beamlet prototype laser demonstrated that the NIF frequency converter design is well understood and expected to meet conversion efficiency requirements. Early experiments by Van Wonterghem et al. tested the basic NIF converter scheme and achieved third-harmonic conversion efficiencies of 80%.¹ These experiments were performed with the crystals held in separate mounts and at atmospheric pressure. Subsequent experiments with NIF prototypical converter hardware in vacuum yielded lower efficiencies that were traced to vacuum-related degradation of the anti-reflection (AR) coatings.² As shown in Figure 4-4, conversion efficiencies for 3ω energy of 73% were achieved under NIF-like conditions, which included the use of rapidly grown crystals. The Beamlet data for a 1.5-ns square pulse is compared to a full-model result that includes, in addition to basic crystal properties, the actual pulse shape, beam diffraction and depolarization, conservative estimates of two-photon absorption, SBS-suppression bandwidth, and measured crystal birefringence variations. The full model predicted 75% conversion in the “as-operated” condition and 80% with normal AR coatings. These results were instrumental in validating our frequency conversion predictive capability and give high confidence that the desired 80% peak-power efficiency can be achieved with stabilized coatings. Progress in stabilizing coatings in vacuum is described elsewhere.³




Figure 4-4. Measured Beamlet frequency conversion efficiency showing good agreement between data and the full model.

A significant investment has been made in instrumentation and techniques to obtain and ensure precise angular alignment of the crystals, which is essential for efficient operation. The angular tolerances for fabrication, assembly, and tuning of the frequency converter are driven by the overall system design, in which the crystals and focus lens are mounted together in a single compact assembly called a Final Optics Cell (FOC), with the input surface of the lens serving as a mechanically rigid alignment fiducial for the high-aspect ratio crystals. In operation, the reflection from the lens is viewed with an off-axis sensor in the NIF Transport Spatial Filter (TSF), while the reflections from the crystals are dumped into the walls of the beam transport tubes at the interface between the switchyard and target bay. As a result, the converter is designed to operate with the 1ω beam incident on the focus lens at a specific angle of 0.58 mrad, and with the doubler and tripler tipped at 12.580 and 10.580 mrad relative to the beam, respectively. To simplify cell fabrication and maintain compact component spacing for ghost focus control, the FOC components are all tipped in the same plane, such that the doubler is tipped about its extraordinary axis (insensitive direction) and the tripler about its ordinary axis (sensitive direction). Successful implementation of this design thus requires tight angular tolerances in three areas: how well the crystals are cut, how well they are registered to the lens, and how well the lens return is aligned in the TSF.

Crystal cut is monitored during fabrication with the aid of two instruments: a full-aperture phase-stepping interferometer and a small-beam pulsed laser metrology apparatus called the Crystal Alignment Test System (CATS). The interferometer measures transmitted wavefront through the crystal at two orthogonal polarizations to quantify its spatial variations in birefringence, which have been found to be good predictors of spatial variations in phasematch angle.² The result is a map of the expected distribution of phase match angles in the plate, relative to the mean. To determine the mean, the CATS measures the absolute phase match angle relative to a standard at a finite number of locations across the crystal. The orientation of the standard is known precisely to within $\pm 10 \mu\text{rad}$ and $\pm 15 \mu\text{rad}$ for doublers and triplers respectively (internal angle, 3σ), as a result of a careful series of frequency conversion measurements in the Optical Sciences Laser Facility in which temperature was controlled to $\pm 0.05^\circ\text{C}$, wavelength was controlled to $< 0.05 \text{ \AA}$, and angle was measured with sub arc-second accuracy. The localized, absolute data from the CATS is then correlated with the full-aperture, relative data from the interferometer to determine the mean orientation of the plate, with an estimated uncertainty of $\pm 15 \mu\text{rad}$ and $\pm 30 \mu\text{rad}$ for doublers and triplers respectively (internal angle, 3σ). Additional details on the use of interferometry and CATS to orient crystals are provided in Section 4.1.3.

Registration of the crystals to the lens is accomplished in the final optics cell, a monolithic aluminum structure that provides full-perimeter support for the optics via a set of precision-machined mounting lands. The lands are flat to $4 \mu\text{m}$ and oriented precisely to each other to within $\pm 7 \mu\text{rad}$ (internal angle, 3σ). In operation the crystals and lens are cleaned and coated and inserted into the cell, where they are held in place against the lands with a compliant force of approximately one pound per inch. Extensive testing on the Beamlet prototype laser demonstrated that this load was sufficient to hold the components in place and maintain a crystal surface figure of better than $5 \mu\text{m}$ P-V. The assembled cell is then installed in the integrated optics module, where actuators provide angular adjustment capability about two orthogonal axes to tip the cell and align the lens reflection to the sensor in the TSF.

Early estimates placed the uncertainty in aligning the lens reflection in the TSF as high as $40 \mu\text{rad}$.⁴ However, subsequent calculations that assumed realistic figure errors for the lens, transport mirrors, and TSF optics, and the specified mechanical tolerances for the hardware in the TSF, indicate that the actual error will be on the order of $10 \mu\text{rad}$ or less.⁵ Using this latter estimate, a worst-case roll-up of all of the alignment errors yields an uncertainty of 30 to $45 \mu\text{rad}$ (internal angle, 3σ) for our ability to align the crystals. For comparison, the frequency conversion error budget specifies an allowable alignment error of $10 \mu\text{rad}$ 1σ , based on early assumptions that the converter fabrication and assembly errors would be measured to this accuracy prior to installation on the NIF.² In fact, the converter tolerances are such that in many cases, metrology of the assembled unit will not be required. The exception will be for experiments requiring precise power balance in the foot of high-contrast ignition-type pulses, where the conversion efficiency is the most sensitive to alignment errors of the angularly biased doubler. The device specified to perform this function is the Converter Assembly Verification Equipment, or CAVE, which will simulate on-line low-power rod-shot

tuning by (1) measuring 3ω conversion efficiency as a function of FOC tilt angle using a full-aperture pulsed beam of a few hundred kW/cm^2 and (2) quantifying the optimum angle of the FOC using a NIF-like alignment system, including fiber-optic alignment source, pinhole assembly, sensor, and a telephoto lens system to simulate L4. Detailed design and implementation of the CAVE is planned for late in the NIF commissioning cycle.

References

1. B. Van Wonterghem et al., "Performance of a prototype for a large-aperture multipass Nd:glass laser for inertial confinement fusion," *Appl. Opt.* **36**, 4932 (1997).
2. P. Wegner et al., "Frequency converter development for the National Ignition Facility," UCRL-JC-129725, Oct. 30, 1998.
3. Hunt, J. T. et al., *National Ignition Facility Performance Review—1999*, UCRL-ID-138120-99, "Final Optics Assembly, Optics and Coatings," September 1999, pp. 4-31–4-45.
4. E. Bliss, "NIF alignment system accuracy for frequency conversion crystals," NIF-0008425, April 24, 1998.
5. P. Wegner et al., "Recommendation to reintroduce a wedged lens in the FOA," NIF-0039912, Dec. 9, 1999.

4.1.1 Frequency Conversion Crystals; Growth and Finishing

Manufacturing of the KDP crystals in the size and volume required for the NIF was a significant challenge. Doubler, tripler, and switch crystals of large ($41\text{ cm} \times 41\text{ cm}$) size are needed with high optical quality and low bulk and surface damage thresholds. Conventional growth methods would take up to two years for each crystal, and this threatened to limit the NIF schedule. Several years ago an investment was made in a new rapid-growth technique for the NIF crystals. In the past three years, significant progress, surpassing expectations, has been made in achieving rapid growth of high-quality crystals by acquiring scientific understanding of the growth process and subsequently engineering reliable hardware. To date, the pilot production runs have produced full-sized boules that potentially yield up to 25% of the initial system requirement on schedule. Figure 4-5a shows a KDP crystal in a rapid-growth tank nearing growth to its full size. Figure 4-5b shows a similar KD*P large crystal also nearing the end of its growth cycle. This success has basically removed crystal growth from the list of significant project issues. Further details on the rapid-growth process are given in Section 6.2.



Figure 4-5a. KDP crystal in rapid-growth tank.




Figure 4-5b. KD*P crystal in rapid-growth tank.

The crystals need to be finished to exacting specifications of flatness and transmitted wavefront. Specifically, they need to be flat to better than five 5 micrometers and be smooth on scale lengths of 120 μm . LLNL has supported development of finishing machines for use by our vendors. Five overall machining steps are required, with the final two processing steps demanding the most significant hardware development. To accomplish this, a prototype semi-finishing machine has been built by LLNL personnel and tested to cut full-size KDP to better than the required 5- μm flatness. A commercial tool manufacturer is being selected to build a production version of this system. A final finishing machine that is anticipated to meet specifications has been built by Moore Tool Company and is undergoing acceptance testing. Crystal finishing is now in a good position for meeting the full NIF requirements. Figure 4-6 shows a full-size doubler crystal being prepared for testing in the Optics Processing Lab.




Figure 4-6. Full-size doubler crystal being prepared for testing.

4.1.2 Finishing Development

Producing finished crystal optics with an aspect ratio of greater than 40:1 from large boules poses a variety of fabrication challenges. During the 1994–96 period, it was demonstrated that it is possible to finish crystals to the specifications needed to achieve acceptable beam quality for the NIF. The challenge over the past few years has been to develop a set of machines and processes capable of meeting those specifications with

cost-effective finishing. Substantial progress was made during 1999 towards this end, with acceptance testing of a new band saw and plate profiling mill and a demonstration that the boule-facing machine and prototype semifinishing machine could both meet their surface finish and flatness specifications. Acceptance testing of the final finishing machine is proceeding with encouraging results.

The crystal finishing specifications are written in terms of rms surface errors in various bands of spatial wavelengths. These specifications are summarized in Table 4-1.

Table 4-1. Crystal finishing specifications written in terms of surface errors.

Spatial wavelengths	Rms surface roughness	Rms transmitted wavefront	Rms transmitted wavefront gradient
0.12 – 0.0025 mm	3 nm	–	–
0.5 – 0.12 mm	2 nm	–	–
0.5 – 6.0 mm	4.2 nm	–	–
6 – 33 mm	–	6.4 nm	–
33 → mm	–	–	11 nm/cm

The semi-finishing machine was among the most challenging of the new machines.¹ In our design (see Figure 4-7), three z-axis slides provide approximately 100 mm of vertical travel and up to 0.25 degrees of tip and tilt of the workpiece platform. The wide range of travel is needed to accommodate different thicknesses of material and various orientation corrections that will need to be made at different points in the blank fabrication process. Two test cuts each were made on an aluminum block (dry) and a z-cut slab of KDP (with oil on the part). The results, summarized in Table 4-2, show that the machine can meet the specifications. The KDP result is especially significant because it is the material of interest.

Table 4-2. Metrology results for dry aluminum and wet KDP cutting tests on the prototype flatness machine.

Property	Specification	1 st cut	2 nd cut
Aluminum (dry)			
Figure @ 633 nm	$2.5\lambda^*$	1.2λ	1.6λ
Roughness-avg (nm-rms)	12	10	10
Angle adjustment (μ rad)	5	0.3	n/a
KDP (wet)			
Figure @ 633 nm	$2.5\lambda^*$	1.9λ	1.0λ
Roughness-avg (nm-rms)	12	7.8	6.5

*one-half the specification for finished crystals



Figure 4-7. Aluminum block on the prototype semi-finishing machine after a dry cut test. The aluminum flycutter is suspended vertically just behind the test piece.

The new final finishing machine was designed and fabricated by the Moore Tool Company (see Figure 4-8). The workpiece is held in a vertical plane by a vacuum chuck mounted on a hydrostatic bearing slideway. A lightweighted flycutter mounted to an air-bearing spindle is attached to a second orthogonal hydrostatic slideway that positions the flycutter to a desired depth of cut. The workpiece slideway feeds the part beneath flycutter to skim-cut the crystal to achieve final surface roughness, waviness, and wedge. Initial acceptance testing during the early months of 1999 found that the machine did not meet many of the performance specifications. Moore subsequently made a significant number of modifications to the machine, and it is currently undergoing additional acceptance testing. Preliminary metrology on the first test cut on a copper-coated aluminum block indicated that the machine is performing close to specifications, but further testing and analysis is required for a firm conclusion.



Figure 4-8. The Moore Final Finishing machine has now been assembled and is undergoing acceptance testing. The reflection of the Moore nameplate is visible on the copper test-cut part mounted to the workpiece slideway.

4.1.3 Crystal Fabrication Alignment

Perhaps the most difficult specification is producing finished crystals oriented with respect to the crystallographic axes such that the average phase-matching angle for frequency conversion is accurate to $\pm 15 \mu\text{rad}$ (external angle) for second-harmonic generators and $\pm 30 \mu\text{rad}$ for third-harmonic generators (THG). Since the phase-matching angle is strongly dependent upon use temperature and wavelength, extremely precise finishing machines and metrology tools are required for process control during crystal fabrication. Furthermore, variations in optical properties across the crystal mean that the optimum fabrication with respect to the crystallographic axes must be determined by averaging many points over the crystal clear aperture.

Variations in frequency conversion efficiency across the crystal correlate very well with variations in crystal birefringence determined simply by subtracting the transmitted wavefront through crystals at two orthogonal polarizations corresponding to the “o” and “e” transmission axes. This technique, now called “orthogonal polarization interferometry” (OPI), provides a valuable alternative to mechanically complex two-dimensional scanning for determining the relative phase-matching angles

across the full aperture of crystals. However, OPI determines only the relative values across the crystal, not the absolute phase-matching angle needed for NIF production.

The Crystal Alignment Test System (CATS) design relies on OPI to determine the distribution of phase-matching angles across a crystal. Small-beam frequency conversion data from the CATS will measure absolute phase-matching angles for a line of points along the crystal. This data is then correlated with the OPI data to provide a map of the absolute phase-matching angle at all points on the crystal and an average phase-matching angle that is used to correct crystal orientation during final finishing operations.

The CATS design, shown schematically in Figure 4-9, is conceptually very similar to the existing methods for determining an absolute phase-matching angle in that each measurement is referenced to a precisely oriented proof crystal. Two parallel beams pass through the proof and test crystals. Pentaprisms, which are mounted on a single diamond-turned chuck, are used to produce two parallel beam paths, thereby greatly simplifying the mechanical and controls requirements of the system. Because frequency conversion crystals are insensitive to out-of-plane rotation errors in the pentaprisms, there is no need to compensate for slide straightness errors or adjust beam alignment at each point in the scan line. During operation, one beam samples the proof, while the other samples the crystal to be tested. The mount is rocked, and frequency-converted power is measured in both channels. The resulting tuning curves are curve-fit to determine the phase-matching angle offset between the two crystals.

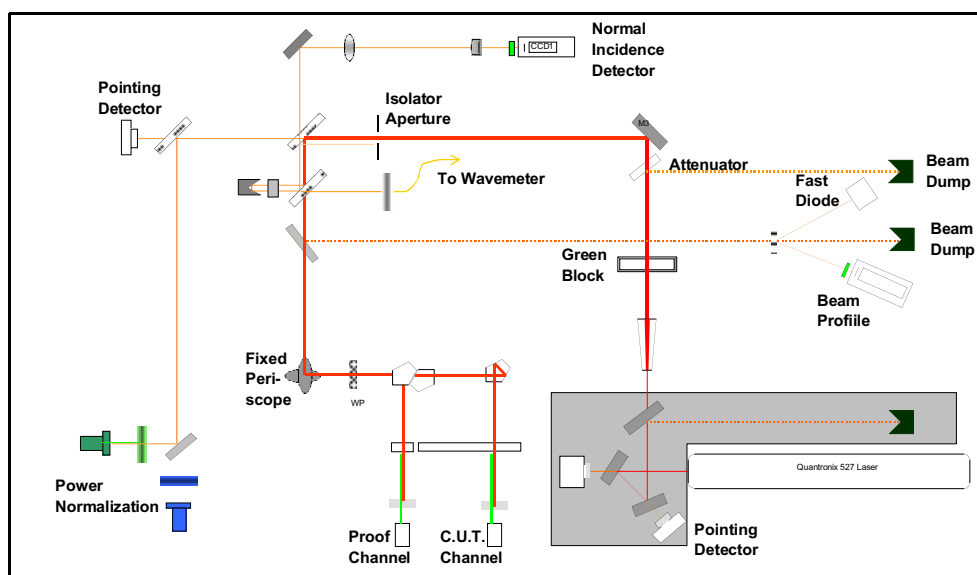


Figure 4-9. Schematic layout of Crystal Alignment Test System for measuring the phase-matching angle of a test crystal relative to a precisely calibrated proof crystal.

The use of calibrated proof crystals in CATS greatly reduces the sensitivity of the measurements to systematic errors, including temperature, laser wavelength, pointing jitter, etc. For example, by comparing crystals to a known proof, only the relative temperature between the proof and the crystal being tested must be controlled. The system is also insensitive to variations in laser wavelength between the NIF-use wavelength and the CATS laser, since both the proof and crystal being tested see the same wavelength. There is only a minor error introduced if the deuteration level of a THG crystal being measured is different than the THG proof.

In summary, the Moore Final Finishing Machine, semi-finishing machines, and CATS, together form the primary foundation needed to meet the most difficult of the crystal finishing specifications. The Moore Final Finishing Machine is undergoing final testing. Mechanical assembly of the Prototype Semi-Finishing Machine is complete, and test cuts on aluminum and KDP met specifications. Acceptance tests are mostly completed on other more standard machines for earlier fabrication steps. A new CATS design takes advantage of new metrology techniques, and construction is nearing completion.

References

1. R. C. Montesanti, et al., "Diamond Flycutting Machine for Producing Flat Surfaces on Half-Meter Scale Diamond-Turnable Optics," presented to the American Society for Precision Engineering, Monterey, CA, November 2, 1999.
2. P. J. Wegner, et al. "Frequency Converter Development for the National Ignition Facility," Proc. 3rd Intl. Conf on Solid State Lasers for ICF, Monterey, CA, June 7–12, 1998, *SPIE*, Volume **3492**, Part One (1999).

4.2 Lifetime of Optics in Operation

A major issue for the reliable and cost effective operation of the NIF is to have optics meet lifetime goals in the range of 500 to 1000 shots or more when operated with average fluences of up to 8 J/cm². Lifetime is somewhat arbitrarily defined as the limit at which obscuration reduces the transmission by 3%. The most significant limit to optics lifetime is the initiation and subsequent growth of damage leading to the loss of transmission through scattering, reflection, and absorption. In some of the final Beamlet experiments completed in 1998, at fluences in the range of 6 to 8 J/cm² significant optical damage was observed on the KD*P tripler, the silica lens, and the diffractive optics. This initiated damage grew to major size in approximately 25 shots. A strong focus of our effort is to understand and mitigate these sources of damage initiation and the causes for damage growth.

Damage generated in NIF optics can in general be described by two mechanisms; that associated with the initiation of damage and that associated with the subsequent

growth of this damage. As shown in Figure 4-10, initial damage sites are typically small, on the order of 20 μm to 50 μm in diameter. If damage was initiated at only a small number of sites, and if this damage did not appreciably grow, the optical transmission would remain high. However, depending on the laser fluence, damaged area can grow relatively rapidly and create significant obscuration after a relatively small number of shots. Consequently, we have focused effort into two activities, one directed toward reducing the number of initiation sites and the other to understanding and controlling the growth of initiated damage. The approaches to understanding damage, including our recent success in reducing initiators and our work in understanding and controlling the growth of damage, are described in subsequent sections.

Our basic model for surface damage initiation centers on the hypothesis that initiation occurs at small (1- μm -diameter for fused silica and ~ 10 - μm -diameter for DKDP) local absorbers that exhibit high relative absorption in the ultraviolet (UV) ($\sim 500\text{ cm}^{-1}$). A local absorber may be an impurity from the manufacturing or finishing process or an embedded particle from handling. The requirement of a high absorption coefficient may be reduced somewhat if the structural defects cause local electric field enhancement and therefore local increases in laser intensity at the initiation site. The energy of an incident laser pulse is highly absorbed at this site and creates a high-temperature and high-pressure plasma. The high temperature and consequent shock wave from the high pressure create an initiated damage site that we believe is characterized by a combination of broken chemical bonds, local stoichiometric changes to the surface material, and fracture growth.




Figure 4-10. Example of NIF initial damage sites.

Our basic model for damage growth follows from the initiation process. It is assumed that the initiated material creates increased absorption at the 3ω wavelength such that a subsequent laser pulse will produce an even greater absorption, further

damage, and thus a sustained damage growth process. This growth continues with each pulse, covering an ever-larger area of the optic.

Our approach to mitigating the UV optical damage thus centers on a two-fold effort to reduce the number of damage initiation sites and to significantly slow the rate of damage growth. We have initially focused on studying uncoated fused silica (SiO_2) and DKDP as the most expedient way to study the basic phenomena. Once we understand and have control of this situation, we will add to our evaluation the additional complication of antireflective coatings and environmental contamination.

4.2.1 Reduction of Fused Silica Damage Initiation

In the process of manufacturing and finishing optics, there are conditions and situations in which contaminants enter the process, resulting in sites that can subsequently cause damage within the bulk or on the surface of the optic. We believe that these contaminants are of nm to μm in size, and once initiated, they produce damage of 20 to 50 μm in size. Our work to reduce damage initiation has focused on identifying contaminants and structural flaws that correlate with laser damage sites. With this information, the steps in the manufacturing process where these defects are introduced can be identified, and the manufacturing process can be modified to prevent their introduction. During the past several years, significant progress has been made in reducing the number of sites per optic at which damage initiates. Figure 4-11 shows the most recent result of a lens made of Corning glass and finished by SVG-Tinsley Corporation. As can be seen, for fluences up to 5 J/cm^2 , the number of damage initiation sites for pure surface-initiated damage has been reduced to the instrument detection limit of less than $0.004/\text{cm}^2$. Two full-scale lenses were damage tested. One resulted in no surface-initiated sites, and the other resulted in only a single site. There were approximately 10 sites per optic initiated at the surface by bulk inclusions. We plan to eliminate these by using a better grade of starting material. Also clearly observable from the figure is that the reduction in damage initiators is primarily associated with fluences up to 5 J/cm^2 . Above this value, the density of initiators climbs back to the earlier values. This behavior suggests that the recent improvement was associated with a low fluence initiator, and there is a yet to be identified contributor responsible for the initiation at higher fluence. We are developing and evaluating more sensitive detection techniques to identify this next level of initiator. Once the initiator is identified, we fully expect to be able to find where it is introduced into the manufacturing process and eliminate it.




Figure 4-11. Example of most recent result of a lens made of Corning glass and finished by SVG-Tinsley Corporation.

One initiator diagnostic technique being evaluated is photoluminescence (PL) imaging of the polished surfaces. Figure 4-12 shows an example of a PL image of a polished surface before and after laser damage initiation. There is a correlation between the location of many of the damage sites with pre-existing luminescing defects. The increased luminescence may be due to contaminants or electronic defects in the silica or to locally increased electric fields at mechanical flaws. Other approaches to initiator identification include analysis of the plume of material ejected during the damage process, using optical and mass spectroscopies. This approach will be useful if the initiators are predominantly absorbing contaminant particles.




Figure 4-12. Example of a PL image of a polished surface before and after laser damage initiation.

4.2.2 Understanding and Reduction of Fused Silica Damage Growth

With the initiated sites in SiO_2 being in the range of 20 to 50 μm in diameter, damage initiation itself does not produce major beam obscuration once the number of initiators is reduced to a few thousand per optic. Thus, at first evaluation, damage initiation is not a problem. However, experiments show that the damaged areas grow in size with subsequent shots at a rate that is approximately exponential. This rapid growth of the damage leads to a large obscuration fraction even for a small number of initiating sites. Thus, unless we can produce optics with only a few damage initiation sites, it is critical for long optics lifetime that we be able to control damage growth.

We have begun to systematically characterize damage growth in SiO_2 . Figure 4-13 shows a typical damage growth curve at $8 \text{ J}/\text{cm}^2$ where a spot initiated at 225- μm size grows to over 1-mm size after 30 shots. Figure 4-14 shows that this damage has a threshold for growth at about $5 \text{ J}/\text{cm}^2$ and a growth coefficient that is fluence dependent. Data was taken in the Slab Lab at approximately 10-ns pulse duration and in the Optical Sciences Lab (OSL) at a range of pulse durations from 1 to 10 ns. There appears to be only a very small change in growth rate due to pulse duration. The observed threshold for growth appears to be stable in that growth at a site can be started and stopped by moving above or below threshold.




Figure 4-13. A typical damage growth curve at 8 J/cm² where a spot initiated at 225-μm size grows to over 1-mm size after 30 shots.




Figure 4-14. Example showing that typical damage has a threshold for growth at about 5 J/cm² and a growth coefficient that is fluence dependent.

We believe that the damage growth in fused silica is closely linked to modification of the silica during damage initiation. The high-temperature plasma associated with the initiation or growth phenomena and its associated shock wave in the material creates structural and chemical defects in the SiO_2 that can enhance absorption of subsequent laser pulses. Several experimental techniques have indicated laser-induced modification of the silica. Photoluminescence spectroscopy indicates the formation of absorbing defects within the damage sites. In Figure 4-15, showing the photoluminescence spectra measured at three locations within a damage site, three distinct luminescence peaks are apparent. It may be reasonable to attribute the peaks to oxygen deficiencies (440 nm), broken Si-O bonds (650 m), and nonstoichiometric material, possibly Si-clusters (560 nm). Photothermal microscopy, which detects laser-induced temperature rises, has also indicated increases in absorption within damage sites. Indications of densification of the silica is provided by Raman spectroscopy, which detected the collapse of the open-ring configurations (i.e., six-member rings or larger), common in the three-dimensional random network of silica, in favor of small-ring configurations (3- and 4-member rings) associated with high-density structures. Densified material can increase local absorption by reducing the silica bandgap or by increasing the laser-induced electric field. Changes in the chemical structure on the laser-damaged silica was also indicated by x-ray photoelectron spectroscopy. Changes in the emission spectra show a partial reduction of the silicon valance state from +4, expected for SiO_2 , to +3, indicating broken bonds and the possible loss of oxygen from the material. Based on these experimental results, insight into the complex nature of the laser-modified silica is emerging. Continuing studies are focussed on quantitatively correlating the observed changes in the silica material with the measured damage growth rate and on determining the thicknesses of the various modifications needed to develop approaches to mitigation of the growth process.

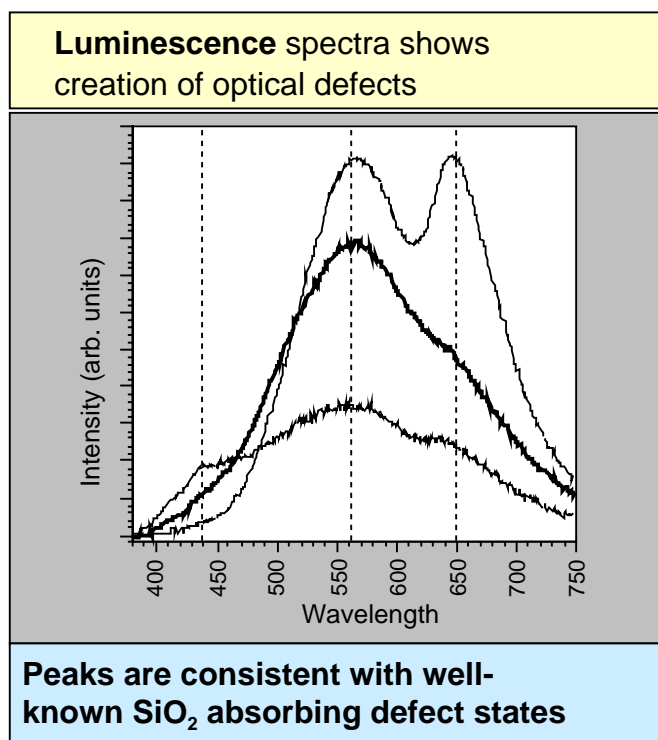


Figure 4-15. Example showing the photoluminescence spectra measured at three locations within a damage site, with three distinct luminescence peaks apparent.

4.2.3 Mitigating Damage Initiation and Growth

We are working along multiple paths to solve the high-fluence damage problem. Along one path, we would reduce the density of initiation spots to a level at which there would be a reasonable probability of having less than one initiator per optic. We could then scan optics with a high-fluence laser before use and, if no damage sites occurred, accept that optic for use. Optics that have damage initiators would either undergo a form of processing to make the site benign, that is, mitigate the damage growth, or be sent back to the manufacturer for reprocessing.

Along the second path, we are developing techniques to stop or slow the damage growth. Toward this goal we are working to better understand the chemistry and physics that determines the growth. If the growth is indeed due to modified material, it is possible that this material could be removed over just the small initiated site, leaving pure, damage-resistant fused silica at the damage site. This approach could work if the number of initiators can be made reasonably small, on the order of 10 to 100 per optic. For higher concentrations of initiation sites a more global processing such as thermal treatment or etching of the full optic may be needed.

Figure 4-16 shows a block flow diagram of a scheme to process optics for high-damage performance. As shown, as optics arrive from fabrication, they are sent to an off-line laser scanning system to initiate potential damage sites. Optics are rejected if

there is an excess number (amount to be determined) of sites. Once damage is initiated, optics are sent to a mitigation station such as a local ion etcher, where the damage at individual sites is removed. The optics are rescanned to qualify and then sent to the field for operation. During operation, optics are routinely inspected for damage, and if it occurs, they are returned to the mitigation station for reprocessing. With this ability to arrest damage effectively developed, it will become highly advantageous to detect any newly initiated damage soon after it occurs. With early detection and changeout of optics, long operating lifetimes can be expected.

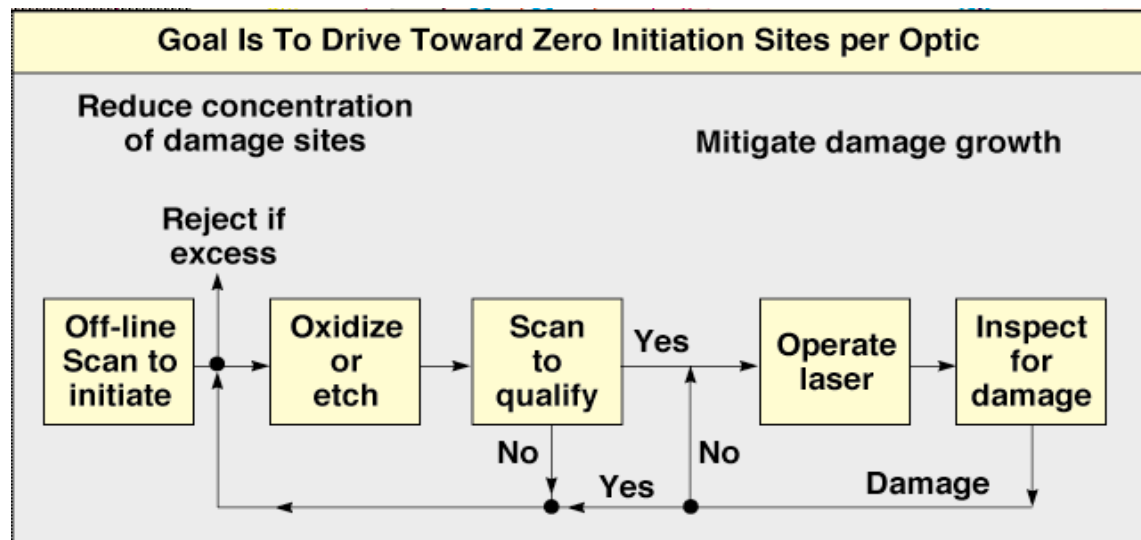


Figure 4-16. A block flow diagram scheme for processing optics for high-damage performance.

A third approach to growth mitigation involves using modified silica compositions that would minimize the formation of absorbing material during the initial damage pulse. Changes in the hydrogen, water, and fluorine content of silica has been reported to influence the laser-damage resistance of silica used for deep UV lithography applications. We are working with the silica vendors to develop and test modified-silica compositions for application at 351 nm.

Significant improvement has already been achieved during the past three years. Figure 4-17 shows the projected optics lifetime parametrically displayed against the number of expected damage initiators per optic and the damage growth rate. The systematic improvement in expected lifetime for the NIF at 1 MJ versus laser fluence primarily resulted from reduction in the number of initiator sites. We expect to achieve additional improvement as we also begin to work on mitigating the damage growth rate.




Figure 4-17. The projected optics lifetime parametrically displayed against the number of expected damage initiators per optic and the damage growth rate.

4.2.4 Identification and Elimination of Mechanisms Leading to UV Damage of DKDP

The design of UV laser systems, including the final optics package for the NIF laser, is limited to a large extent by the laser damage performance of the third-harmonic generator, which is made from deuterated potassium dihydrogen phosphate (DKDP).^{1,2} The damage may occur either in the bulk of the material or on the surface. In the past, the surface damage threshold of KDP and DKDP crystals was usually higher than the bulk damage threshold for diamond-turned surfaces. However, as the bulk damage threshold has increased over time, primarily in response to better filtration of growth solutions and purer salts, surface damage has become more of a concern.

Bulk damage. Damage morphology is strongly suggestive that light-absorbing initiators deposited throughout the crystal during growth “explode” when exposed to sufficient fluence, thereby forming a melted and fractured zone several micrometers in diameter. Generally, these bulk damage sites do not grow significantly at the fluence levels of interest to the NIF, but the number of sites has to be low enough to prevent unacceptable levels of scattered light on subsequent laser pulses.

The traditional way of measuring bulk damage thresholds is to test dozens to hundreds of spots with a small (~1-mm Gaussian) laser beam. If the fluence at each spot is ramped up until damage occurs, it is called an R/1 test. If the fluence at each spot is a fixed fluence, it is called an S/1 test. The R/1 damage probability curve occurs at higher fluence because the low-fluence shots precondition the material by dissipating the damage sites, as shown in Figure 4-18 for a z-cut sample from a rapid-growth boule (BD6) grown in a 1000-L tank. Experiments in prior years established a correlation between the traditional damage probability curve measured for a 7.6-ns pulse and the amount of obscuration allowed by the NIF specifications for a 3-ns pulse. The boundary between acceptable and unacceptable material is represented by the R/1 profile of rapid-growth sample DKDP11, also shown in Figure 4-18.

The interpretation of the acceptability of boule BD6 is complicated by a discovery in late 1999 that the damage probability curve depends on whether the sample is a z-cut (parallel to the crystal base) or tripler cut (rotated 59° from the base). If that were not so, the S/1 curve in Figure 4-18 would indicate that BD6 material could be used directly at 8 J/cm² at 350 nm without preconditioning. However, results in Figure 4-19 for a similar boule show that the damage curve of a z-cut is nearly two times higher than a tripler cut of the same material. The result was confirmed by exposure of the same samples to a 1-cm-diam flattop beam on the Optical Sciences Laser. As a result, preconditioning of BD6-type material would be required to operate satisfactorily at full-NIF fluence.




Figure 4-18. Comparison of the S/1 and R/1 damage probability curves of sample BD6 with the R/1 curve for sample DKDP11, which represents the boundary between acceptable and unacceptable material. The shaded fluence ranges represent NIF beam modulation ranges after scaling for differences in pulse length assuming a $\tau^{0.5}$ dependence.




Figure 4-19. S/1 damage probability curves for various samples taken from 1000-L tank DKDP boule BD7. The tripler orientation samples have a systematically lower damage threshold than the z-cut samples for reasons that are not yet understood.

We investigated in parallel the feasibility of off-line 3ω conditioning^{3,4} and the possibility of growing material that would not need preconditioning. Results from these experiments indicate that either approach would work. For the former, preconditioning, we raster-conditioned several crystals with both 350-nm light and 308-nm light from an excimer laser. Results given in Figure 4-20 show that either the Optical Sciences Laser or the excimer laser can improve the S/1 damage distribution, especially the initial 20% that is most important for meeting NIF specifications to close to the R/1 damage distribution curve. To be most effective, the raster conditioning should be performed in two or three steps of increased fluence. The excimer laser would probably be more cost effective, given its higher average power output for comparable cost.

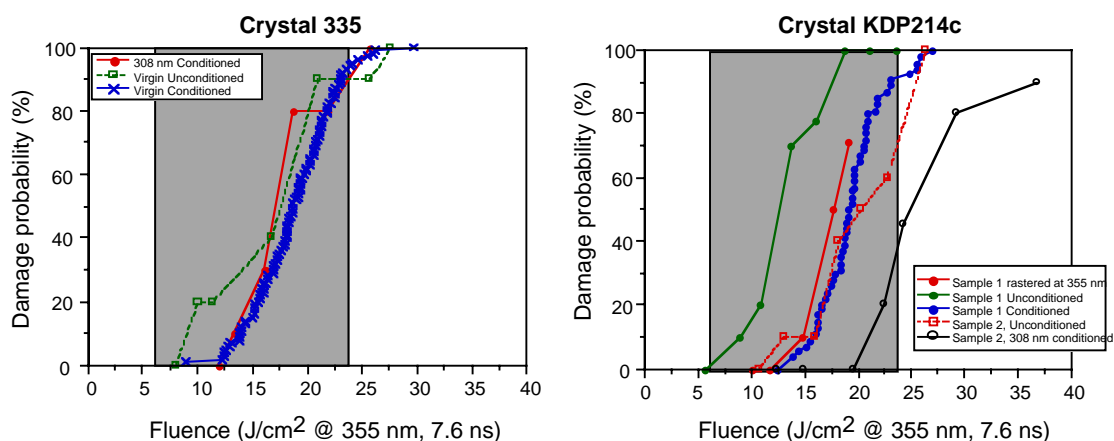


Figure 4-20. Improvement in S/1 damage probability curve with preconditioning at 355 (left) and 308 (right) nm. It appears possible to improve the on-line S/1 damage curve to the R/1 damage curve.

To determine whether we could grow material not requiring preconditioning, we grew a set of 18 DKDP samples to systematically explore the relationship between 3ω bulk damage properties and the range of plausible engineering parameters: production salts of varying purity, growth temperature, resaturation sequence, and growth tank material. All samples were grown by reducing the temperature as the crystal increased in size, so that the growth rate remained approximately constant as the dissolved salt concentration decreased. Of the parameters explored, the most significant appears to be initial growth temperature. Twelve crystals grown from approximately 65 to 45°C had the midpoint of their damage probability curve range from 17–22 J/cm², while the corresponding range for six crystals grown from 45 to 25°C was from 25–40 J/cm²—all z-cuts. The reason for this difference is not known, but it appears possible to implement on an engineering scale. Results for the crystal having the highest damage resistance are shown in Figure 4-21. Even allowing a twofold reduction for a tripler cut, this material would tolerate the full NIF power at 8 J/cm² without preconditioning.




Figure 4-21. Bulk damage probability distribution for a z-cut of the best rapid-growth DKDP sample. Even after allowing for a twofold lower value for a tripler cut, the material could be installed on the NIF and operated at full fluence without preconditioning.

Surface damage initiation. Some experiments on Beamlet had unacceptable levels of surface damage on the output of the tripler. There are several possible initiation mechanisms for surface damage, including near-surface bulk damage, imbedded particles from the finishing process, and surface contamination. An example of near-surface bulk damage is shown in Figure 4-22. The inverted square pyramid damage shape is characteristic of the orientation of the crystal as cut from the boule. In addition, near-surface particles in excess of that expected from near-surface bulk inclusions have been detected by both photoluminescent imaging and photothermal microscopy. These techniques will be used in future experiments to help identify and eliminate, as much as possible, the most important factors for surface damage initiation.




Figure 4-22. Near-surface bulk damage for a sol-gel coated DKDP crystal. The areal density is proportional to the bulk damage density times the effective depth for surface eruption. For a NIF-quality crystal, the areal density is predicted to be about 0.5 damage site/cm² for an 8-ns, 15-J/cm² pulse.

As for fused silica, DKDP surface damage growth appears to require a threshold fluence. Below that fluence, damage sites are stable for thousands of shots, and significantly above that fluence, they grow to centimeter size in a matter of a few shots. Near the threshold, however, the likelihood that the growth will become catastrophic appears to be random. In other words, two damage craters of similar size exposed to the same fluence may follow the same size trajectory for dozens or more shots, whereupon, one will grow catastrophically, as shown in Figures 4-23 and 4-24. The reason for this is not understood, but it appears to be related to the presence of cracks instead of a more melted morphology. Further experiments are planned to better understand the mechanism of surface damage growth.

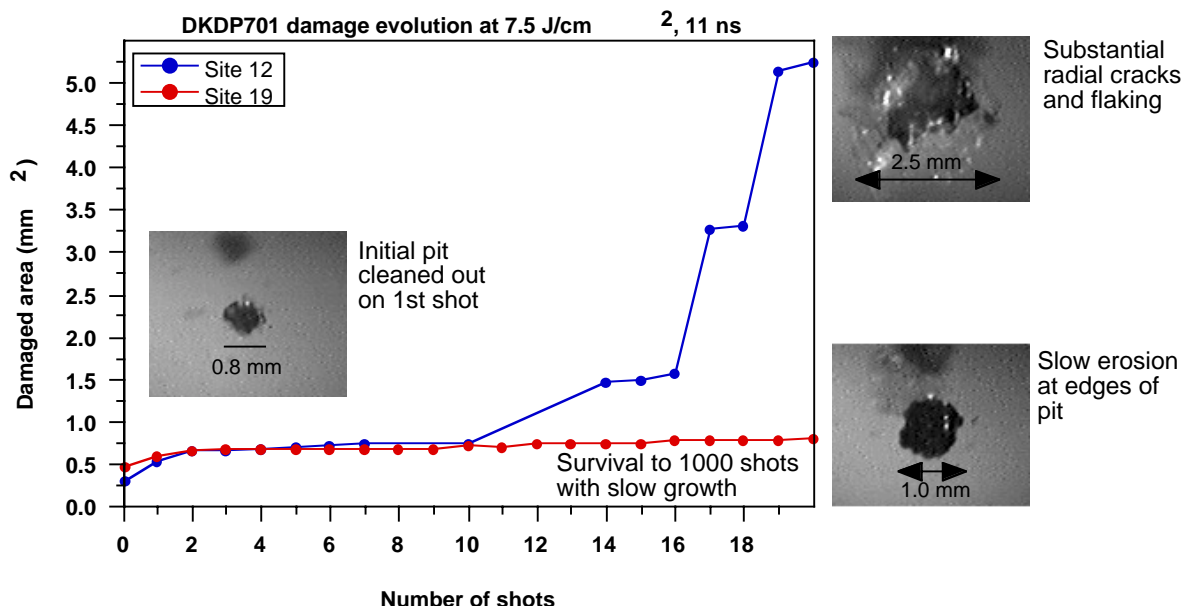


Figure 4-23. Measurements of surface damage in bare KD*P show a break point in growth at 3ω fluences above 7 J/cm² (11 ns). The transition to catastrophic growth involves a change in damage morphology, but the mechanism is unknown.




Figure 4-24. Damage growth proceeds rapidly at fluences above 8 J/cm² (10 ns), but sometimes has an induction time. LL3 site 4 was much smaller than other sites (diam = 36 μ m) and was laser induced at ~ 30 J/cm² (7.6 ns). LL1 sites were produced by mechanical indentation. The LL3 site survived prior exposure to 1000 shots at 5 J/cm² and 7.5 J/cm² (10 ns).

If DKDP initiation sites cannot be completely eliminated, mitigation of predamaged sites may be as simple as etching with water. Defect formations at DKDP damage sites absorb at 351 nm and fluoresce with a peak at 560 nm. Recent experiments have shown that the absorbing material from damage sites can be removed using a dual-orifice nozzle (water injected in the center and withdrawn in a surrounding annulus—see Figure 4-25). Damage tests on the stability of the etched crater will first be conducted. Other more sophisticated etching techniques, such as ion milling or deep UV ablation, may be tried if the etched craters are not adequately stable at NIF UV fluences.




Figure 4-25. UV-induced fluorescence from an original DKDP damage site (left) and a damage site dissolved by a dual-orifice nozzle (right). The maximum intensity is 50 times higher in the original damage site.

References

1. A. K. Burnham, et al., "Producing KDP and DKDP Crystals for the NIF Laser," *ICF Quarterly Report*, LLNL Report UCRL-LR-105821-99-2.
2. N. P. Zaitseva, J. J. De Yoreo, M. R. Dehaven, R. L. Vital, K. E. Montgomery, M. Richardson, and L. J. Atherton, "Rapid growth of large-scale (40-55 cm) KH_2PO_4 crystals," *J. Crystal Growth*, **180**, 255 (1997).
3. J. Swain, S. Stokowski, D. Milam, F. Rainer, "Improving the Bulk Laser Damage Resistance Of Potassium Dihydrogen Phosphate Crystals By Pulsed Laser Irradiation," *Appl. Phys. Lett.*, **40**, 350 (1982)
4. S. G. Demos, M. Staggs A. Burnham, J. J. DeYoreo, and H. B. Radousky, "Damage induced material modification in the bulk KDP crystals," *XXXI Annual Symposium on Optical Materials for High-power Lasers*, Boulder, Colorado, October 4–7, 1999.

4.3 Optics and Coatings

4.3.1 Improved Anti-Reflection Coatings for the Conversion Crystals

For many years, quarterwave porous-silica antireflection (AR) coatings prepared by a sol-gel process have been used on the transmissive optical components of high-powered fusion lasers. Although the optical performance of these sol-gel coatings is no better than coatings prepared by other methods, their laser damage threshold is two to three times higher.¹ It is thought that this enhanced damage threshold results from the extremely high purity obtained by distilling the starting materials and the lack of collateral damage when small particulate contaminants are ejected. Optical component damage has always been the limiting factor for high-power output, so sol-gel coatings have made a significant contribution to high-power laser development. The general issues related to manufacturing sol-gel coatings for the NIF are given in Section 6.4.

Due to their porosity, sol-gel coatings are susceptible to capillary condensation of both organic contamination and moisture into their porosity. Fused silica optical coatings can be made less susceptible to organic contamination by treatment with ammonia and hexamethyl disilazane, as discussed in Section 6.4.5. Further reduction of the organic contamination problem by the use of carbon getters is discussed in section 6.4.5.3. Crystal coatings are more difficult, and the special issues related to crystal coatings are discussed in the following sections.

Single-layer, porous, colloidal-sol coatings can meet the optical and damage requirements for all of the NIF's transmissive optics, except for the two mixed-wavelength surfaces on the frequency conversion crystals. Development of two-layer broadband coatings and coating processes that are compatible with the conversion crystals' requirements has been one area of focus for improving NIF laser AR coatings. The second area of active research is the long-term performance of both single-layer and broadband AR coatings using porous sol layers. The sol-gel AR coatings consist of a layer of porous, near-spherical silica particles, 10 to 40 nm in diameter, randomly stacked on the substrate surface. Interparticle voids combine with the particles' interior porosity to produce a film refractive index very near to the square root of the indices of many common optical substrates (1.46 to 1.52); nearly 100% transmission can therefore be obtained. Unfortunately, such highly porous coatings have high specific surface areas and are therefore very susceptible to contamination by vapor adsorption from their environment. Furthermore, water vapor adsorption by the sol coating results in etching of the surfaces of KDP crystals, which in turn generates wide-angle scatter and loss of energy on target. Methods to reduce water vapor uptake or to protect the reactive KDP surface have yielded notable improvements in the lifetime performance of porous sol-coated KDP crystals.

4.3.2 KDP Anti-Reflection Coating Development

The KDP/DKDP conversion crystals are the only optics in the NIF laser chain that must simultaneously transmit mixed wavelengths of light with very high efficiency. This requirement has implications for selecting acceptable AR-coating designs and thin-

film processes. To independently optimize the coating for each crystal surface, it has been necessary to develop a spin-coat process for 41-cm-square optics with acceptable uniformity, repeatability, and low particulate contamination. In addition, for the output surface of the doubling conversion crystals and the input surface of the tripling conversion crystals, the single AR-coating design must be simultaneously optimized for both 1053-nm and 525-nm light. This can be accomplished by coating the surface with an AR layer having an optical thickness that produces a minimum of reflectance between these two wavelengths. Figure 4-26 compares the calculated reflectance from such a single-layer “compromise” coating (currently being used on deuterated tripling crystals) with that of a two-layer broadband AR coating used on KDP crystals for Nova and Beamlet.




Figure 4-26. The calculated reflectance from a single-layer compromise coating compared with a two-layer broadband AR coating used on Nova and Beamlet crystals.

The single-layer coating surpasses the 0.7% goal for reflection losses established for these surfaces, while the two-layer coating meets the NIF's requirements.

4.3.2.1 Spin-Coat Process Development

The sensitivity of compromise coatings to variations in optical thickness increases the need to maximize their uniformity. Unfortunately, the spin-coat process used for Nova and Beamlet crystals exhibited relatively poor uniformity and repeatability when compared to dip coating. The Nova/Beamlet spin coater relied on fairings to minimize air turbulence at the edges of the square crystals.

A commercial spin coater that was developed for coating large, rectangular substrates (such as flat-panel displays) was purchased for NIF KDP optics. In this

system, the cover spins with the optic, creating a quiescent, highly saturated air pocket with low air turbulence and slow solvent evaporation. Nevertheless, judicious selection of spin coat solvents is still required to achieve the desired within-part and part-to-part thickness variation. For the Nova-Beamlet era colloidal sols, coating uniformity across a 37-cm part was improved from a standard deviation of 0.1%, to one of 0.04% by replacing two-thirds of the ethanol solvent with secondary butanol. Solvent selection for the experimental HMDS-treated sols (discussed in more detail below) is still under way. Table 4-3 compares the capability of our current spin-coating process to the dip-coat and spin-coat processes used for Beamlet.

Table 4-3. A comparison of 1064-nm reflectance uniformity for 37-cm crystals that were coated by various processes.

Process	Solvent	$\sigma(\%R)$
Beamlet - Dip Coat	Ethanol	0.08
Beamlet - Spin Coat Over Dip Coat	2-butanol over ethanol	0.15
NIF - Spin Coat	Ethanol	0.1
NIF - Spin Coat	2-butanol:ethanol (2:1)	0.04

4.3.2.2 Broadband coating development

The two-layer KDP coating used on the 27-cm Nova and the 32-cm Beamlet crystals consists of a methylsilicone bottom layer (with a refractive index of 1.43) that has been overcoated with porous silica. The silicone coating is GR-650,TM a commercial material available from Techneglas. It is applied from an alcohol solution and requires curing at 140°C. After being heat-cured, the material is insoluble and will not wash off when the colloidal sol coat is subsequently applied from a second alcohol solution.

The first 37-cm KDP crystal was coated with a two-layer broadband coating on one surface and a single-layer, 3 ω sol coating on the opposing surface using the new (NIF) spin coater. Despite initial concerns that, with such a large crystal, either shrinkage of the silicone layer or the thermal cycle itself might distort the reflected wavefront, there was no measured change in the reflected wavefront. Table 4-4 compares the NIF's goals to the measured reflectance and uniformity of this coated crystal.

Table 4-4. The performance of a 37-cm KDP crystal, coated by the NIF process and the spin coater.

Coating Design	Wavelength (nm)	Reflectance (%)	σ (%R)
2-Layer Compromise	1064	0.47	0.02
2-Layer Compromise	532	0.34	0.07
Single-Layer 3 ω	352	0.09	0.02

4.3.2.3 Improved broadband coating development

The NIF's frequency-tripling crystals must be deuterated to suppress parasitic oscillations from stimulated Raman scattering. Deuterated crystals have catastrophically failed when heated above 100 to 120°C, so they are not compatible with the thermal treatment required by the existing broadband coating. Two routes have been proposed to eliminate the need for elevated cure temperatures. One route relies upon modifying the chemistry of the intermediate-index layer to allow it to "cure" without elevated temperature. The other route is to replace the alcohol solvent in the colloidal silica with a fluid that is not a solvent for the GR650 resin. The first route produces an intermediate-index layer that cross-links upon standing in air to become insoluble in ethanol. The other route relies upon the modification of the colloidal silica particles via HMDS-treatment to produce a sol that is dispersible in heptane or similar nonpolar solvent.

The commercial GR650 silicone resin that we have used for many years is a polymer prepared by an acid-catalyzed hydrolysis of very high purity methyltriethoxysilane (MTS). The material produced by this reaction is incompletely cured and can be applied from an alcohol solution. This coating, as deposited, requires additional heat treatment to produce a fully cross-linked insoluble product. However, the polymerization of MTS, a tri-functional monomer, can be enhanced by adding a tetra-functional monomer, such as tetraethylorthosilicate (TEOS). The extra functionality speeds up polymerization and increases the cross-linking required for rapid insolubility. The copolymerization of TEOS and MTS can be carried out in any proportion to obtain a soluble prepolymer that can be cured to insolubility with little or no additional heating, depending on the relative proportions of the monomers. A coating made from a 50/50 mixture of MTS and TEOS will become insoluble in ethanol after standing overnight at room temperature.

The refractive index of the final product ranges between 1.42 and 1.43, which is indistinguishable from the original GR-650 silicone and produces the same optical performance. The nominal laser damage threshold of this low-temperature-cure

silicone is equivalent to that obtained with the commercial GR-650™ silicone and exceeds the expected peak fluences of the combined 1053- and 525-nm laser light.

4.3.3 Degradation of the Optical Performance of Sol-Coated KDP

Figure 4-27 shows the transmittance degradation of the rapid-growth frequency-doubling crystal, RG8B-2, which was used for the final conversion and high-damage-threshold campaigns in the Beamlet vacuum mule. This optic is representative of the doubling and tripling crystals being operated in vacuum. Unlike the fused-silica optics, the decrease in KDP/DKDP transmittance is not accompanied by a commensurate increase in reflectance, and the original transmittance is not regained when the optic is washed with a suitable solvent.




Figure 4-27. Transmission loss over time for the Beamlet rapid-growth doubler crystal, RG8B-2.

This transmission loss is attributed to the formation of discrete etch-pits on the crystal surface under the porous sol coating. These pits, generally one-to-two microns wide and a few hundred nanometers deep, have characteristic geometries that depend on the crystal surface's orientation. Figure 4-28 shows the diamond-shaped pits that are found on Type-I crystals, such as doubler RG8B-2. The Type-II crystals, used for third-harmonic generation on Beamlet, develop triangular or trapezoidal pits. The Z-cut crystals, used for the Pockels switch, develop square etch pits. The etch pits on the Beamlet's tripling crystals elongated to become "slits" up to 40-microns in length, which cause wide-angle scattering. This scatter loss dominated the crystal's performance, far

outweighing any transmittance drop from organic contaminants adsorbed into the porous sol.




Figure 4-28. Characteristic pit shapes found on KDP and DKDP crystals: (a) diamond-shaped pits that are found on Type-I crystals, such as doubler RG8B-2, (b) triangular and trapezoidal pits found on Type II (tripling) crystals, and (c) square pits found on z-cut (switch) crystals.

4.3.4 Formation of Etch Pits

Examination of crystals from the historical database dating back to the early 1990s revealed that only crystals that had a sol coating and were exposed to humidity formed etch pits. Controlled experiments confirmed that both the porous sol coating and humidity are required for development of etch pits.

Further examination of these crystals suggests that the nucleation and/or growth processes are diffusion, rather than reaction rate, limited. As shown in Figure 4-29, the etch pit shape is self-similar regardless of size; this suggests that the dissolution and rearrangement of KDP ad-atoms on the pit surface is fast enough to maintain an equilibrium shape as the pits grow. A second indication that diffusion is controlling etch pit nucleation and/or growth rate is the development of depletion zones around sites of preferential nucleation (scratches), as shown in Figure 4-30. The width of the depletion zone increases with increasing relative humidity, suggesting that the diffusion rate of KDP through the porous sol increases as the amount of adsorbed water increases. Furthermore, both the size of the individual etch pits and total volume of etch pits increases with humidity, as seen in Table 4-5.

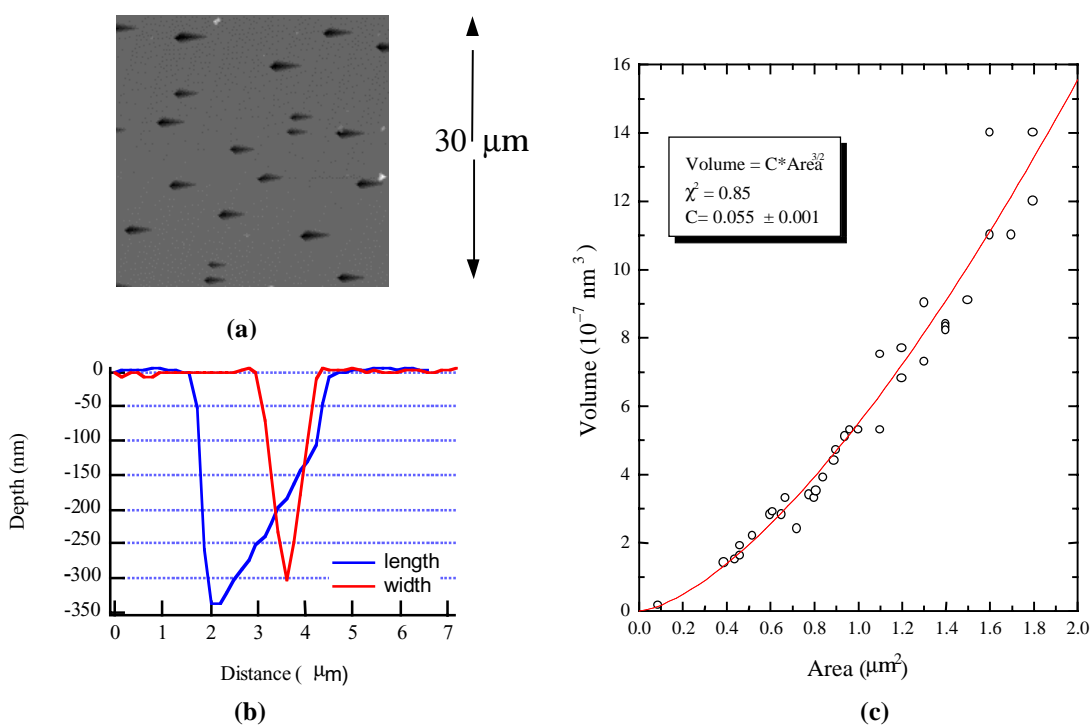


Figure 4-29. The etch pits are self-similar; knowledge of any characteristic pit dimension defines the pit volume and area. (a) AFM image of etch pits on doubler crystal after laser ablating the coating, (b) profile of etch pit along two axes, and (c) correlation between volume and area of etch pits over a range of sizes.

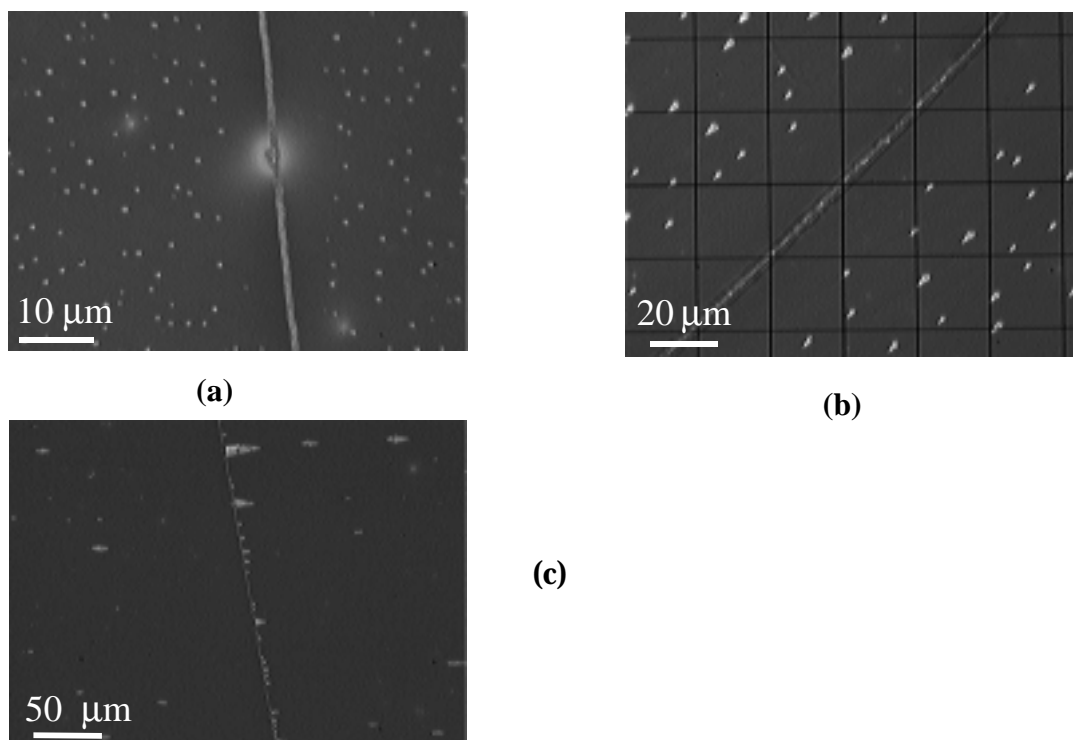


Figure 4-30. Development of denuded zones around scratches at (a) < 5%, (b) 55%, and (c) 75% RH.

Table 4-5. Etch pit size, density and total volume and denuded zone length for type II rapid growth KDP crystals as a function of relative humidity. Data for crystals shown in Figure 4-30.

Relative Humidity (%RH)	Denuded zone (μm)	Number density (per cm^2)	Average length (μm)	Total etch pit volume (μm^3 per cm^2)
< 5%	~6	7×10^6	0.5	3×10^3
55%	~20	6.2×10^5	4.4	2×10^5
75%	~50	7.3×10^4	2.5 or 15	4×10^5

Engineering Model

These experimental observations lead to the following model for etch pit formation. Water from the environment is adsorbed into the porous sol coating. Dissolved KDP is wicked up into the sol as etch pits nucleate and grow. Growth continues until the water reservoir in the sol is saturated with KDP. Hence, the total etch pit volume should reach an equilibrium value, which is a function of the ambient relative humidity (RH).

An engineering model was constructed based on the premise that two primary processes dominate etch pit formation. The model proposes that very small pits are initiated in the KDP surface as a result of the adsorbed moisture in the coating at a rate of formation R_N

$$R_N = k_N (C - C_s)^m \quad (1)$$

followed by a pit growth process which is assumed to be controlled by the ability of KDP to radially diffuse from the pit into the adsorbed moisture in the surrounding sol coating at a growth rate R_G

$$R_G = k_G (C - C_s) \quad (2)$$

where C is the local concentration of KDP in the water reservoir, and C_s is the concentration of KDP at saturation. The nucleation coefficient, k_N , contains the material and surface-state dependent detachment energies. The mass transfer coefficient, k_G , is a function of the effective diffusivity and pit spacing. The existence of depleted zones led to the assumption that initiation is a strong function of undersaturation. Hence, the nucleation rate is represented by a power law dependence on undersaturation with $m > 1$.

This simple model is consistent with observed trends in pit morphology when relative humidity is varied. Increased relative humidity is expected to result in higher diffusivity in response to the higher water saturation in the sol coating. As shown in Figure 4-31, the model predicts that the time scale to form pits will decrease, and the size of the pits will increase proportional to the square root of the change in diffusivity. In addition, the model predicts that the number of pits per unit area will decrease proportional to the square root of the change in diffusivity.




Figure 4-31. Calculation of etch pit growth rate (size) as a function of the growth coefficient, k_G

To gain further insight into the behavior of etch pits in a wide variety of environments, the model was reduced to a series of dimensionless equations. The resulting dimensional analysis allows us to draw a number of conclusions about how well the model predicts our experimental observations. The model is qualitatively consistent with the experimental observations for change in etch pit size, number density, and growth rate as a function of changes in relative humidity. The data are insufficient to confirm the model prediction that the ratio of the time scale to the number of pits should be invariant with saturation.⁶ Furthermore, the model predicts that pit size and total pit volume should increase with sol coating thickness, while the number density should be invariant, when all other parameters are held constant. Instead, our experimental observation is that pit size is invariant (within our measurement error), and pit density, and hence volume, increase with increasing coating thickness. More recently we have obtained incontrovertible evidence of coarsening at 75% RH. We are currently refining this first rudimentary model to include coarsening behavior and to obtain better concurrence with the experimental observations.

Etch Pit Mitigation

The proposed model has formed the basis of a 5-pronged etch pit mitigation strategy:

1. Apply a barrier coating to reduce or eliminate water transport to the KDP surface.

2. Modify the KDP surface to increase the energy barrier to nucleation.
3. Modify the porous sol surface chemistry and/or structure to reduce adsorption of water or transport of KDP through the sol.
4. Eliminate the porous sol coating by utilizing a fully dense, low refractive index material.
5. Eliminate exposure to humidity.

4.3.5 Barrier Coatings

A loss of transmission, attributed to a surface roughening of contaminated KDP surfaces in the presence of water vapor, was observed sporadically on Nova.⁴ In the early 90s, we began using a 122-nm-thick silicone layer on all Nova crystal surfaces as an environmental, as well as an optical, coating. When Nova was decommissioned last summer, conversion crystals that had been on-line since as early as 1990 were examined for etch pits. Eighteen crystals installed in 1995 had no silicone coating, and all 18 exhibited slit-like etch pits characteristic of type II KDP. Of the remaining 63 crystals with silicone coating, only 3 had small etch pits. The remaining 60 crystals, including 18 crystals that were in ambient humidity service on Nova for over nine years, had no etch pits. For NIF, the doubler coatings will combine the benefits of a 160°C thermal anneal (see below) and a protective GR650 barrier coating layer to eliminate etch pits.

Unfortunately, silicone coatings were not used on the 37-cm Beamlet KDP crystals due to concerns about wavefront distortion and cannot be used with DKDP crystals due to the requirement of a thermal cure. In our laboratory, the more brittle room-temperature-cured silicone coating has been observed to reduce, but not consistently eliminate etch-pit formation on DKDP. In addition, none of the silicone coatings exhibit sufficient 3ω damage threshold to be used on the output surface of the DKDP tripling crystal. High molecular weight commercial polymethylmethacrylate (PMMA) also proved a successful barrier to etch pit formation, but had insufficient 3ω damage threshold to be used for NIF triplers.

4.3.5.1 KDP Surface Modification

Two approaches have been explored for passivating the surface of KDP crystals to eliminate etching: thermal annealing the crystal after diamond turning and direct chemical passivation. Once again, thermal anneal is only applicable for KDP, not DKDP. Thermal anneal has proven to retard, if not eliminate, the formation of etch pits upon exposure to humid environments. In addition, thermal annealed surfaces form many very small etch pits that exhibit less scattering than the larger etch pits formed on equivalently handled unannealed crystals (Figures 4-32 and 4-33).⁵

Chemical passivation holds the promise of eliminating etch pits on both KDP and DKDP crystals. A commercial zirconate coupling agent (KZ-55, Kenrich Petrochemicals, Inc.) was tested both as an additive in the final finishing (diamond-turning) process and as a post-treatment. Efforts to incorporate the compound into the diamond turning oil proved futile as the compound readily reacts with both water and KDP, precipitating zirconium-rich compounds onto the diamond turning machine plumbing and leaving behind a blotchy, insoluble organic layer on the KDP. Reaction

of a monolayer of zirconium alkoxide by dipping the optic in a freshly diluted solution of the coupling agent in toluene provided robust protection against etch pits; however, the application process must be carefully controlled to avoid a deleterious excess of the coupling agent. Work is currently focused on several other promising chemical passivation agents that may afford the protection of the KZ-55 coupling agent without the deleterious side effects and extreme process sensitivity.




Figure 4-32. Comparison of scattering for unannealed and annealed Type I KDP crystals and a crystal with a polymer barrier coating beneath standard sol gel.




Figure 4-33. Example of (a) thermal annealed surface that formed many very small pits exhibiting less scattering than the (b) larger pits formed on equivalently handled unannealed Type I crystals.

4.3.5.2 Modification of the Porous Sol Surface Chemistry and Structure

Hydrophobic silica sols have been prepared by reacting the polar surface silanol groups with hexamethyldisilazane (HMDS) in an ethanol solution, as shown in Figure 4-34. The excess HMDS can then be evaporated, along with the ammonia catalyst, before the sol is applied to the crystal surfaces. DKDP crystals coated with these sols showed no etch pits after exposure to 55% RH for several months.⁵ As shown in Figure 4-35, water adsorption isotherms confirm that these coatings adsorb less water than untreated coatings at equivalent humidity and fail to exhibit capillary condensation at any humidity. The lack of visible etch pit formation for crystals coated with HMDS-modified porous sol coatings is consistent with the hypothesis that the equilibrium etch pit volume (i.e., scatter) is controlled by the amount of water adsorbed into the coatings.




Figure 4-34. Reaction of polar surface silanol (Si-OH) groups with hexamethyldisilazane (HMDS) to produce hydrophobic HMDS-treated sols.




Figure 4-35. Water adsorption isotherm for HMDS-treated sol vs. untreated sol as a function of relative humidity. The ordinate is the fraction of the void that is filled with water.

The nonpolar trimethyl siloxyl surface groups that populate the HMDS-treated porous sol have an additional advantage; they permit dispersion of the colloidal sol in non-polar solvents such as heptane. Since the methyl silicone optical coating is insoluble in heptane, there is no need for a thermal cure to allow a porous sol overlayer to be applied. Hence, the same methyl silicone polymer that is used for the doubler broadband coating can be used for the tripler.

Initial attempts to synthesize HMDS-treated sols resulted in foaming, flocculation, and gellation of sols during the solvent exchange, upon standing or during coating. Recent efforts have improved the synthesis process robustness and efforts are now turning to solvent selection to enable uniform spin-coat application and demonstration of adequate 3ω damage performance.

4.3.5.3 Eliminate Porous Sol Coatings

Etch pit formation can be eliminated by removing either the porous sol coating (the water reservoir) or the water (relative humidity). Both methods have proven successful. A 3ω coating utilizing AF2400™, a low-index, high-purity fluoropolymers developed for the photonics industry,^{7,8} performed without etch pit development when exposed to 55% relative humidity for a month. Similarly, KDP optics coated with untreated porous silica sol did not develop etch pits when held in low (<1% RH) humidity. The AF2400 coating has several drawbacks: (1) higher index than porous sols results in higher Fresnel loss, especially for 1–2 ω coatings, (2) adhesion and 3ω damage threshold are unproven for large optics, and (3) neutron degradation is expected to limit useful optical life when NIF achieves full operating fluences. Hence, development of

these coatings has low priority. Quantifying the kinetics of etch pit growth as a function of RH for the various proposed sol and crystal surface treatments remains a high priority in developing a predictive model to guide the processing, storage, and operations environments for KDP crystals. The option of controlling exposure to humidity remains a viable, if costly, option for elimination of etch pits when a silicone barrier coating is not applicable.

4.3.6 Coating Conclusions

AR coating designs and materials have been developed to meet NIF requirements for KDP/DKDP optics. Intermediate index materials, such as a room-temperature-cured silicone have been developed that allow broadband coatings to be applied to deuterated tripling crystals to allow optimal performance at two wavelengths. Substantial progress has been made toward a spin-coating capability, which will allow us to place an optimum-thickness coating on each surface of the conversion crystals. Efforts have shifted from manufacturing process development to developing methods to better control and minimize the interaction of these coatings with their use environments. Demonstration of the effectiveness of barrier coatings, modification of the KDP surface, and modification of the colloidal-sol surface chemistry provide multiple pathways to eliminate etch pits when sol-coated crystals are exposed to ambient humidity.

4.4 Summary—NIF Final Optics Assembly and 3ω Optics

The key elements of the NIF 3ω optics are defined according to function and required performance. Critical experiments have been completed on the Beamlet system demonstrating that the required 80% peak frequency performance can be achieved. Development of rapid-growth technology for KDP and DKDP has been very successful, with full production capability now in hand. A new generation of fabrication and finishing machines has been developed, and we are successfully completing demonstrations of their performance. Except for a final problem associated with the generation of etch-pits in DKDP, we can successfully produce the required antireflective coatings and maintain them at high quality. The etch pit problem is understood, and solutions are being evaluated and developed. The focus of our efforts is on improving the expected operational lifetime of the 3ω optical components under high fluence loading. We are studying the phenomena that initiate optical damage and influence its growth, determining causes and developing approaches to reduce initiators and to reduce the growth rate or delay the onset of growth. We believe that we have demonstrated the required performance of NIF 3ω optics and that achieving the desired lifetime can follow.

References

1. I. M. Thomas, "High Damage Threshold Porous Silica Antireflective Coatings," *Appl. Opt.* **25**, 1481 (1986).
2. I. M. Thomas, et al., "Processes for the Elimination of Fogging on KDP Crystals Prior to and During Use in Laser Systems," *Boulder Damage Symposium*, Boulder, Colo., October 23–25, 1991 (UCRL-JC-108273).
3. I. M. Thomas, "A Two Layer Broadband Antireflective Coating Prepared from a Methyl Silicone and Porous Silica," *Proc. of SPIE Sol-Gel Optics IV*, **3136**, 215–219 (1997).
4. M. R. Kozlowski et al., "Influence of Diamond Turning and Surface Cleaning Processes on the Degradation of KDP Crystal Surfaces," *SPIE International Symposium* (July 21–26, 1991) UCRL-JC-105835.
5. E. K. Wheeler, et al, "Scatter loss from environmental degradation of KDP crystals," *Boulder Damage Symposium*, Boulder, Colo., October 23–25, 1999 (UCRL-JC-134819).
6. C. B. Thorsness, "Modeling of KDP Pit Growth," LLNL internal memo, Oct. 28, 1999.
7. I. M. Thomas and J. H. Campbell, "A Novel Perfluorinated AR and Protective Coating for KDP and Other Optical Materials," *Boulder Damage Symposium* (Oct. 24–26, 1990).
8. S. Papernov, et al., "Laser-Induced Damage in Optical Materials," 1991, *SPIE* **1624**, 314–318 (1991).

5.0 WAVEFRONT CONTROL	5-1
5.1 NIF Wavefront Control System.....	5-1
5.2 Deformable Mirror (DM).....	5-2
5.2.1 DM Design and Development	5-2
5.2.2 NIF Baseline Deformable Mirror.....	5-2
5.3 Wavefront Sensor.....	5-13
5.3.1 Hartmann Sensor Design Concept and Validation Status	5-13
5.3.2 Hartmann Sensor Risk Mitigation	5-14
5.4 Wavefront Controller	5-18
5.4.1 Wavefront Controller Design Concept and Validation Status.....	5-18
5.4.2 Wavefront Controller Risk Mitigation.....	5-19
References	5-21

5.0 Wavefront Control

Rich Zacharias, Bruce Woods, Scott Winters, Lewis Van Atta, Ken Jancaitis,
William Rivera, and Tim Arnold

5.1 NIF Wavefront Control System

Each NIF beam includes a wavefront correction system to compensate for aberrations induced by laser thermal effects, optics figure errors, etc. The wavefront correction system consists of an adaptive optic, a wavefront sensor, and a computer control system that work together in closed loop until just prior to the laser system shot to apply the requisite phase compensation. The NIF adaptive optic is a square 400-mm clear-aperture, 39-actuator deformable mirror (DM). The wavefront sensor is a miniature Hartmann sensor that is based on a lenslet array and a charge-coupled device (CCD) camera. It is housed in the Output Sensor Package. The wavefront control computer system consists of an image processor computer that reads the Hartmann sensors and a pair of wavefront control computers that control the actuators of eight DMs via dedicated power amplifiers. System requirements for the NIF Wavefront Control System and its subsystems are described in several published documents.¹⁻⁷

A 100% Review for NIF Wavefront Systems (W.B.S. 1.7.3) was held in November 1998. Reviewer comments (NIF-0018084) were published as part of the NIF Technology Annual Report for 1998 (NIF-0025088 [DRAFT]). See *National Ignition Facility Wavefront Control Development and Manufacturing*, 1999, NIF-0054450,⁸ for the responses to those comments (NIF-0023305) published in March 1999.

Last year's technology report emphasized analysis of NIF wavefront control performance, the development and qualification of DM test facilities, and initial residual error measurements of the baseline DM. This year's report describes candidate subsystem designs that were considered and some of the steps taken to select and to validate the designs, to assure that they can be built to meet performance requirements, and to assure that they can be produced on budget and schedule. Particular emphasis is given to the DM, which is the wavefront control subsystem with the highest cost and risk. The report also includes a description of the process used to develop DM designs, the rationale for the selection of the baseline design, and the process and rationale for selecting a manufacturer.

5.2 Deformable Mirror (DM)

5.2.1 DM Design and Development

Three DM designs were considered for the NIF. These designs were proposed and developed by Raytheon Optical Systems Incorporated (ROSI), by ThermoTrex Corporation (TTC), and by LLNL. The process whereby DM designs were solicited and developed is described in Ref. 8. By the end of FY98, prototypes of all three DM designs had been built and had undergone preliminary tests. The LLNL design had been selected as the baseline design for the NIF, but the other two designs were still being considered pending prototype repairs, upgrades, and further tests. Initial measurements of the LLNL prototype DM can be found in Section 5 of the NIF Technology 1998 Annual Report (NIF-0025088 [DRAFT]). Ref. 8 describes the ROSI prototype, the results of its testing, and some analysis of its advantages, disadvantages, and risks. Similarly, that document describes the TTC DM. Ultimately, the LLNL design was selected as the NIF DM in February 1999. The rationale for this selection is also described in Ref. 8, which describes as well the selection rationale for the manufacturer.

Initial tests of the baseline design prototype, described in the 1998 Technology Annual Report, demonstrated that the design and fabrication approach could yield an acceptable residual error. Given that an earlier LLNL design had showed excessive residual error on Beamlet tests, this was a very important parameter to validate. However, there were other design and manufacturability issues that still needed to be validated. The remainder of this section describes the results of FY99 design validation and risk mitigation analyses and experiments to support the baseline NIF DM design.

5.2.2 NIF Baseline Deformable Mirror

5.2.2.1 Design Requirements and Concept

The NIF DM must meet stringent performance requirements and must operate in a severe environment. Some of the features and requirements of the DM are shown in Table 5-1, and some parameters of the environment within which it must operate are shown in Table 5-2. A particularly stringent goal is that the DM should have less than 0.025 waves of rms residual error between the DM surface and a true flat surface when the mirror is commanded to be flat in closed loop. (Note: NIF system models showed that this specification can be relaxed somewhat without significant loss of NIF focus performance.³) The 2-J/cm² flashlamp fluence is a particularly severe environmental parameter.

A photograph and concept sketch of the NIF DM is shown in Figure 5-1. This DM is the latest of numerous generations of DMs that have been built at LLNL. It employs lead magnesium niobate (PMN) actuators that move a glass faceplate by

pushing against a stiff metal reaction block. The prototype DM faceplate was fabricated by Zygo and coated by Spectra Physics, and the actuators were manufactured by Xinetics.

Table 5-1. Deformable mirror features and requirements.

Features, Goals, and Requirements	
Parameter	Specification
Actuator pattern	Hexagonal with 39 control points
Residual Error	0.025 waves rms surface (closed loop to flat)
Weight	≤ 200 lbs
Size	Within 452 by 472 mm by 185 mm
Clear aperture	> 400 by 400 mm
Coating	Coating reflectivity ≥99.5% Coating: 0.2% ≤ Transmission ≤ 0.5%
Correction stroke	≥ 4 waves (surface)
Actuator linearity	≤ 8%
Actuator hysteresis	≤ 5%
Actuator lifetime	Actuator lifetime ≥ 10 ⁹ cycles
Actuator bandwidth	≥ 100 Hz (open loop)
Actuator replacement	Replaceable actuators
Mechanical resonance	300 Hz

Table 5-2. Deformable mirror environment requirements.

Environment
2 J/cm ² laser pump flashlamp fluence
EMI of 8 gauss and 13 V/m in a 200-μsec pulse just prior to the laser shot
Relative humidity ≤3%
Class 50 cleanliness on the optical surface and class 100 for the assembly




Figure 5-1. Photograph and concept sketch of the NIF DM showing major components.

There are several major features designed into this DM to allow it to meet its fluence, precision, and reliability requirements. First, the actuators are held in constant compression by a disk flexure rather than being directly connected to the faceplate. Actuators directly connected to the faceplate work against neighboring actuators putting some of them in tension. Tension in the actuators can cause microscopic cracks to propagate, leading to latent failures. The disk flexures also protect the actuators from the shear load of the faceplate. Keeping the actuators in compression improves the reliability of the DM.

Second, the faceplate is held to the actuators indirectly by epoxy shear joints between the aluminum-coated DM posts and stainless-steel cups. Flashlamp exposure tests of early design concepts had demonstrated that a simple, unprotected butt joint would be destroyed by NIF's extreme flashlamp fluence. With this design approach, the joint is on the side of the post, and the geometry and total internal reflection limit the fluence seen by the aluminum coating. The aluminum protects the Hysol 9330 epoxy. This actuator connection approach was successfully tested on the previous-generation large-aperture mirror in Beamlet. The shear joint allows the DM to meet its fluence specification.

Third, the faceplate high-reflectance coating is applied prior to DM assembly rather than afterwards (the typical practice for commercial DMs). To survive the high-fluence NIF beams, the DM employs a Hafnia-Silica high-reflectance multilayer dielectric coating. This coating is applied at high temperatures that would destroy the epoxy bonds in the actuator force train if the coating was applied

after DM assembly. The coating allows the DM to meet its beam fluence specification, but unfortunately, the stress applied by the coating can be significant and is a function of humidity and aging. Coating stress can increase DM residual error and is a significant challenge for this design. Approaches to minimizing coating stress-induced distortions are being investigated at the Laboratory for Laser Energetics (LLE) at the University of Rochester.⁹

Fourth, a butt joint is used to attach the mirror cups to the reaction block. Compared to a direct mechanical connection, this approach makes the DM much less susceptible to misalignment-induced moments on the mirror posts that cause front-face distortions. It also allows for improved precision and lower residual error.

5.2.2.2 Baseline Deformable Mirror Risk Mitigation

5.2.2.2.1 Joint Survival and Creep. After residual error, the next highest-risk issue in the DM design was judged to be failure of epoxy joints. The DM is the end mirror in the NIF main laser cavity. It is in close proximity to the laser slab amplifiers and, as a result, will see light from both the laser beam and the amplifier flashlamps. The beam will be efficiently reflected by the DM front-surface coating, but much of the flashlamp light will penetrate the coating and be deposited into the DM. It has been estimated that the DM will be irradiated with 2 J/cm^2 of broadband flashlamp light from the amplifiers over a nominal $360\text{-}\mu\text{s}$ pulsewidth. The DM must survive and continue to perform within specifications after having been exposed to this fluence for thousands of shots.

The material behind the faceplate that is the most sensitive to the flashlamp fluence is the epoxy used to bond DM components. There are two potential failure modes of the epoxy. First, a joint could release its bond to the joined materials. Second, the epoxy in a joint under tension or compression could gradually creep or flow due to either heating or epoxy material degradation, allowing movement of the joined parts. Since there is limited margin in the stroke of the DM actuators, such creep would reduce the effective correction range of the DM, and at some point, would render the DM unserviceable.

Flashlamp light will be partly absorbed by the DM BK-7 faceplate. Most of the energy will pass through it and be deposited onto the components behind the faceplate. The most vulnerable epoxy joint is the shear joint that attaches the cups to the mirror faceplate posts. The cup joint employs Hysol EA9330 epoxy. A coating of aluminum on each mirror post reflects the flashlamp light, preventing it from reaching the epoxy. Above the joint, the cup incorporates a cutout to trap any epoxy that overflows the joint during assembly and keeps direct flashlamp light away from the epoxy.

The tests of the first-generation DM on Beamlet in 1997 showed no epoxy joint failures, leading to the belief that bond failure is not a problem. There were only 12 full-power shots on this Beamlet campaign, so those tests alone provided

insufficient data to assure that the epoxy joints in the DM would not creep over time.

An actuator in the DM will run out of stroke if a number of conditions that use up its 12-wave specified stroke simultaneously occur: (1) The actuator has little or no margin above its stroke specification; (2) a large fraction of the actuator range is used up in DM assembly to pull-out the primarily-spherical aberration induced by the faceplate high-reflectance (HR) coating; (3) the wavefront correction required at the actuator location is near the extreme of the 4-wave correction specification for the DM; (4) nearby actuators are bucking the actuator so that it needs to apply a large force, which tends to reduce the efficiency between the displacement of the actuator itself and the displacement of the faceplate surface; and (5) creep in the flashlamp-exposed shear joint over the life of the DM has reduced the usable stroke by too large an amount. A budget of 1 μm is assigned for epoxy creep.

Aluminum Mirror-post Coating Attenuation of Flashlamp Light. First, a simple calculation was made to assure that the aluminum coating is thick enough to prevent flashlamp light from depositing energy onto the epoxy. The design aluminum thickness is at least 2000 Angstroms. For wavelengths down to the UV, the depth at which a wave penetrating a nonalkali metal such as aluminum falls to $1/e$ of its initial value can be expressed as $d = c/(8\pi\mu\sigma\omega)^{1/2}$, where c is the speed of light, μ is the material permeability, σ is the material conductivity, and ω is the angular frequency. The vast majority of the flashlamp energy falls between 10 and 0.1 μm . The skin depth for these two wavelengths is 26 and 263 Angstroms, and the attenuation of a 2000-Angstrom layer ranges from about 100 dB to 1000 dB. (Since the BK-7 also attenuates wavelengths longer than 2 μm and shorter than 0.3 μm , these attenuation values are conservative.) Thus, we can consider the aluminum coating on the faceplate posts to be opaque.

Flashlamp Exposure Tests. Recently, a series of experiments were conducted to determine if the flashlamp light striking the cup joint epoxy degrades the joint and/or induces creep. The tests were designed to directly measure creep. At the same time, they would allow the observation of any gross damage to the epoxy—the generation of carbon deposits—and the observation of any damage to the aluminum.

Figure 5-2 shows the experimental setup for the flashlamp exposure tests. Samples of the joints were cut from the prototype mirror used on the Beamlet tests. Each joint was placed under a 90-lb. load, the maximum expected for NIF mirrors. This was accomplished by supporting the back side of the mirror faceplate on two stainless-steel rods, which were captive in slots in the mirror support block, and by pulling on the cup joint using a calibrated series of Bellville springs. This distorted the sample's front mirror surface, and any creep could be very accurately detected by looking at the corresponding relaxation in this surface figure with an interferometer.

Null tests were first performed to determine how much, if any, the samples changed over time due to influences other than flashlamp exposure and to determine measurement repeatability. With the test samples unexposed, it was found that over a several week period, less time than a typical experimental series, there was no detectable change in the samples and that the repeatability of the measurement was about 0.08 waves at 1 μm . The measurement error was dominated by diffraction near the edges of the interferograms and by errors in re-positioning the test samples in front of the interferometer.




Figure 5-2. Flashlamp Exposure Experiment Set-up.

The flashlamp exposure tests are conducted at an accelerated rate relative to the NIF, with only a few minutes between shots. This short time does not allow the samples to return to the ambient temperature before the next shot and causes the epoxy's bulk temperature to rise. In NIF, there are at least four and typically eight hours between shots so the joint bulk temperature never rises above ambient. Therefore, it was necessary to determine how much the bulk temperature could rise in the exposure tests without it inducing creep in the joint.

The temperature at which an epoxy changes from an amorphous solid to a liquid due to a bulk temperature rise is defined as its glass transition temperature, T_g . This temperature was measured for samples of the Hysol 9330 epoxy⁸ and found to be between 114° and 118°F. It was decided that the average temperature of the test pieces should be kept well below 100°F to minimize creep due to bulk heating of the epoxy.

Numerous tests were conducted in the flashlamp exposure facility to establish a baseline shot rate and configuration that would not excessively elevate the average temperature of the sample. The resultant configuration included a nitrogen gas cooling line through the sample chamber. The four lamps on the ends of chamber were turned off. These would contribute heat but would not contribute much useful light due the large angles at which their light approached the samples. In this configuration and with a firing rate of one shot every five minutes, the average temperature of the test facility stayed at about 84°F, well below the 100°F limit for the bulk epoxy. Another null test was made in this configuration, with the samples exposed to the indirect, average-temperature rise, but blocked from direct-light exposure. No creep was measured.

The fluence of flashlamp light in the exposure experiments was estimated based on prior measurements made by Erlandson and Truelson.¹⁰ They measured a fluence of $19.6 \pm 0.5 \text{ J/cm}^2$ near the center of an 8-lamp half-shell with Ce-doped, 1.5-cm bore diameter, 112-cm arc lamps with a nominal pulse width of 600 μs . The current experiments used the same number and type of lamps, the same lamp packing fraction, the same explosion fraction and the same pulse-width. Therefore, it is reasonable to estimate that the current experiments are operating at the same fluence as was previously measured. Uncertainty in the actual exposure chamber fluence was estimated to be as much as a factor of two, due primarily to the differences in geometry of the two experiments.

The test samples were placed into a polished stainless-steel test chamber, which was open at the top to allow flashlamp light to enter. Different fluences were obtained by placing stainless-steel screens over the opening. The attenuation for each of the screens was determined by measuring the transmission of the screen using a spectrophotometer, which sampled a 4.5- by 10-mm area of the screens. Transmissions were very flat over the measured range of 300 to 1000 nm.

Flashlamp testing was conducted at fluences of 1.9, 4.9, and 11.8 J/cm^2 . The test samples were exposed to a total of 1220 shots at 1.9 J/cm^2 with no visible degradation or creep of the test samples. The same was true for exposure at 4.9 J/cm^2 with a total of 1100 shots. After 100 shots at 11.8 J/cm^2 , however, one sample showed noticeable creep, and two of the samples showed visible carbon deposits on epoxy that covered the post outside of the cup. In sample preparation, some of the aluminum coating on the post had been etched away and was not protecting the epoxy as well as it will on NIF mirrors. The epoxy was also not as well shielded from direct flashlamp light, as it will be on the NIF mirrors. After 600 shots at the 11.8 J/cm^2 fluence, all three samples tested showed measurable creep, with the change in surface figure ranging between 0.3 and 0.4 waves at 1 μm .

Figure 5-3 shows a typical before-and-after difference interferograms of joint samples. The processed interferograms were masked and have tilt removed. Notice that the center is higher than the edges for (a), which is what would be expected if

the sample had relaxed due to creep. The sample in (b) has exhibited no measurable creep after flashlamp exposure.




Figure 5-3. Surface difference measurement of test sample. Rise at the center of the sample (a) is attributed to epoxy creep, whereas the sample in (b) showed little creep.

As a (planned) final test, the samples were subjected to 5000 shots at a fluence of 4.9 J/cm^2 , which is about 20% of the expected number of shots that the mirror will see on the NIF. Unfortunately, during the test, the attenuation screen that limited the fluence onto the samples was damaged by the severe flashlamp light, and it came apart, allowing portions of the test samples to be exposed directly to the flashlamps (see Figure 5-4). The average temperature of the test samples when this happened was 90°F with peak temperatures of about 93°F . When the test samples were measured in the interferometer, the creep was found to be between 0.35 and $0.5 \lambda\text{s}$. Much of this creep could have been caused by the bulk temperature rise of the samples during the test rather than by prompt flashlamp heating. An additional test was performed exposing the samples to 3300 shots at a fluence of 6.7 J/cm^2 . The screen survived this test intact. The test showed minimal creep with changes in surface figure ranging from 0.08 to $0.13 \lambda\text{s}$. The steady-state temperature during testing was 86°F with peak temperatures of 89°F .

The tests described above determined a relationship between the creep of the joint and the number of shots. The 3300 shot test at 6.7 J/cm^2 is the most credible experimental series in this nominal fluence level since the attenuation screen stayed intact. Minimal creep was observed during this test. The creep observed on the 5000-shot test at 4.9 J/cm^2 is an upper bound on the NIF creep rate, because there is a possibility that the observed creep was due to bulk heating instead of, or in addition to, prompt heating. The measurements could thus overestimate the prompt-only creep rate that would be seen on the NIF. At the creep rates measured,

the DMs would begin to exceed their 1- μm budget after about 12 to 17 years, and NIF operations might start to see DMs fall below their specified correction range. If this happened, frequent and expensive DM refurbishments would be required. The initial 1000-shot test at 4.9 J/cm² showed no measurable creep, but the measurement uncertainty was such that it would be unreliable to extrapolate from only 1000 shots to the 24,000 shots expected on the NIF.

A single test was conducted to measure creep of the epoxy joints under load when the samples were heated in an oven at 85°F for seven days. These test samples showed creep between 0.19 and 0.25 waves. This data suggests that even at modest elevated temperatures, the epoxy joints may exhibit significant creep. Therefore, DMs should be shipped in insulated containers and storage in non-air-conditioned areas should be avoided.



Figure 5-4. Photograph showing damage to attenuation screen in the flashlamp exposure facility.

No creep was measured in samples exposed to 4.9 J/cm². This fluence allows a margin of a factor of 2.45 over the expected 2 J/cm² NIF fluence. The uncertainty in the expected NIF fluence is estimated to be only about $\pm 10\%$ (see below). The uncertainty in the estimate of the exposure facility fluence (the 4.9 J/cm²) is estimated to be less than a factor of two. Therefore, the factor of 2.5 margin is sufficient to overcome the uncertainty in both the test fluence and the NIF fluence estimate, and thus assures that the DMs will survive. The small creep that was demonstrated in the samples at 11.8 J/cm² shows that the margin is about a factor of five.

Analysis of Flashlamp Fluence at the DM. The predicted fluence at the DM was calculated using a line-of-sight technique.⁸ This approach had been proven to be successful in estimating the flashlamp light leakage into the beamlet optics cavity diagnostic (LLNL internal memorandum LTV-93-018). This model also correctly

predicts the pumping loss (and lower gain) in end amplifier segments as opposed to those in the middle of a chain.

The analysis takes into account light that directly exits the cavity, light that reflects off the disks and disk holders, and light that passes through one or more slabs before it impinges on the DM. In addition, the model tracks not only the visible portion of the output spectrum, which generates the inversion, but also the UV and IR emissions from the flashlamp that are not absorbed by the laser glass. The model however does not take into account reflection loss off the silver and the modifications in the angular distribution of the flashlamp light by the shaped reflectors.

Adding up all of the contributions, we would expect the flashlamp loading of the DM to be 1.9 J/cm^2 at a flashlamp loading (explosion fraction) of 20%, rising to 2.2 J/cm^2 if the explosion fraction is increased to 23%. Based on the assumptions made in the model and the comparisons to the experimental data mentioned above, we estimate the error in these calculations to be about $\pm 10\%$.

5.2.2.2.2 Summary of Modal Analysis Tests. Modal analysis tests were completed on the prototype NIF DM to verify mechanical resonant frequency requirements. The requirement was written to limit the lowest mechanical resonant frequency to 300 Hz. Fundamentally, this requirement is needed to assure successful closed-loop performance of the adaptive optics system and to prevent vibration-induced beam steering. The sample rate (update rate) goal for the control system is 30 Hz, and it is necessary to keep mechanical resonances sufficiently above this frequency. If a mechanical resonance exists within or near the control system bandwidth, then vibrations that coupled to the DM could be amplified and could disrupt the wavefront control system operation.

The LLNL Modal Analysis Lab conducted resonance tests on the prototype DM. The summary report, "SM98-051TGW/MWE,"⁸ overviews the test setup, tests performed, and test results. Of the four mode shapes and frequencies determined for the DM, the lowest was 1100 Hz. This is a factor of 3.6 times the 300-Hz design requirement. This mode describes the twisting of the reaction block. The next lowest mode was the torsion of the mirror about its centerline. These results confirmed that the baseline DM meets the resonant frequency requirement.

5.2.2.2.3 DM Temperature Stability. The BK-7 faceplate has a coefficient of thermal expansion that differs from that of the aluminum reaction block used in the prototype DM. This means that after the two subsystems are joined by epoxy, the glass will be stressed as the temperature differs from that at which they were joined, which will be 68°F , the operating temperature of the NIF amplifier cavity. If the temperature of the DM differs too much from 68°F , the faceplate could shatter. Ref. 8 contains a calculation that shows that with the aluminum block, the faceplate could break with as little as a 10°F difference. It was judged that this represented too great a risk to the DMs on NIF, so an Engineering Change Request (ECR) was submitted and accepted to change to a stainless-steel block. Calculations show that

the faceplate will now survive with a temperature difference of up to 40°F. ECR 649 describes the proposed change and the issues for the NIF.⁸

5.2.2.3 Baseline DM Risk Assessment.

5.2.2.3.1 Advantages and Disadvantages. A major advantage of the LLNL DM design is the fact that the brittle actuators are always held in compression, and they do not support the shear load of the faceplate. This is expected to mitigate the propagation of microscopic cracks that could lead to failure. The LLNL design evolved from earlier generations of DMs produced at LLNL for other programs, which also featured precompression applied to the actuators. While precise failure statistics are not available, clearly these programs demonstrated dozens of DMs of similar complexity, each operating for hundreds of hours with failures in the one-to-few range, depending upon the definition of failure applied. The small Beamlet DM operated for about four years with one failure. This reliability approaches that which will be required of NIF DMs.

A second major advantage of the LLNL DM is that the epoxy joints have been demonstrated to survive the flashlamp fluence they will see on the NIF. This was demonstrated by flashlamp exposure tests (see Section 5.2.2.2.1) and by actual laser shots on Beamlet.

A third advantage of the LLNL DM is the butt-joint that prevents manufacturing errors in the reaction block and flexure assemblies from creating moments on the faceplate posts that add residual error to the assembly. These errors are mitigated by lapping the faceplate cup surfaces and the reaction block flexure surfaces flat, and then marrying the two subassemblies with epoxy.

A disadvantage of the LLNL DM is that it will be difficult to achieve and maintain the required cleanliness. The “nooks and crannies” in the force train assemblies between the faceplate and reaction block are difficult to clean prior to final assembly and nearly impossible to clean after final assembly, and the lapping step will potentially add difficult-to-remove grit to the subassemblies prior to final assembly.

A second disadvantage of the LLNL DM is that the HR coating must be applied prior to final assembly due to the high coating temperature required. The coating is humidity sensitive and stresses the faceplate, which tends to increase residual error.

5.2.2.3.2 Technical/Performance Risk. Successful risk mitigation experiments have reduced the technical and performance risk of the DM to acceptable levels. Most of the remaining risks are associated with cost, production, and yield.

The stainless-steel reaction block design is not proven. The best way to validate the new design is to build one. This will be done as part of DM pilot production.

The redesigned LLNL DM has replaceable actuators, but their replacement without adding an unacceptable increase in residual error has not been validated. Thus, there is a performance risk in the long term (30-year lifetime) at which we would expect to have to replace actuators periodically.

5.2.2.3.3 Schedule Risk. The major schedule risk associated with the LLNL design is based on LLE production of DMs.⁸ The personnel that will be assembling the DMs are of the team that built and maintain OMEGA. If a major problem were to occur with OMEGA, personnel that were fabricating or coating DMs might need to be diverted to address OMEGA issues.

5.2.2.3.4 Cost Risk. It is estimated that the remanufactured LLNL design DMs would cost about \$75K each in production, including BK-7 optics. This cost estimate was based on actual vendor quotes for the pilot production materials on a man-hour estimate based on the actual time to build the prototype LLNL DM and on manpower costs based on LLE labor rates. The original Title II cost allocated to the DMs was \$53K, including faceplate and coating costs. This was later reduced by a Budget Transfer to \$43K after both manufacturer's preprototype ROM production cost estimates came in at around \$25K. (After the prototypes were built and tested, the installed cost for both manufacturer's DMs in production were estimated to be higher than \$53K, when required auxiliary systems were included.)

5.3 Wavefront Sensor

5.3.1 Hartmann Sensor Design Concept and Validation Status

The NIF will employ a miniature Shack-Hartmann sensor to measure the wavefront.^{11,12}

A sketch depicting the operation of the sensor is shown in Figure 5-5. Each lenslet generates a focus spot whose position displacement is directly proportional to the local deviation from total-beam collimation of the portion of the beam that impinges on it. The sensor has been demonstrated to resolve 0.1 wave at 1.053 μm . The prototype sensor uses an array of lenslets manufactured by MEMs Optical Systems Inc. Index-matching fluid is used to adjust the lenslet focal length to the value that provides the required angular sensitivity and range for use on the NIF.

In the NIF, each Output Sensor services two beams. For most diagnostics, the operator selects which beam is to be viewed, but since the Wavefront Control System must operate simultaneously for all beams, two beams are spatially multiplexed onto one sensor.¹³ The sensor monitors a beam picked-off by the diagnostic beam splitter that has been demagnified and is relay-imaged to be somewhat smaller than one-half of the CCD camera array, as shown in Figure 5-5b.




Figure 5-5. (a) Shack-Hartmann sensor concept sketch and (b) focal plane image of dual-beam Hartmann sensor.

By the end of FY98, a prototype Hartmann Sensor (HS) had been built and demonstrated. A design change in the DM (5% scaling of actuator spacing to optimize performance) caused a corresponding change in the spacing of the two array patterns in the HS, but it was judged to be a low-risk minor change; the new NIF design has not been built.

A manufacturing issue that still needs to be resolved is to achieve the correct initial lenslet focal lengths directly or with solid adhesive. The index-adjusting material used to set the sensor focal lengths is a liquid. From a reliability point of view, the use of an index-adjustable adhesive that solidifies would be desirable. It has a more limited index-adjustment range than does the liquid. On the first best-effort manufacturing attempt, the lenslets were specified to have an initial focal length that would have allowed the use of the adhesive, but the manufacturer failed to hit it, so a liquid index match was used for the prototype. Since the liquid index-adjusting material is a back-up option, this is not considered a high-risk issue. The production lenslets will be specified for the desired focal length with best effort, and if the manufacturer fails to meet it, the liquid index match will be used in production.

The remaining risk items associated with the HS all have to do with its performance given the variations of intensity and the limited field of view (FOV) it will encounter in the NIF.

5.3.2 Hartmann Sensor Risk Mitigation

The Hartmann wavefront sensor in conjunction with the video acquisition system and mirror controller must be robust when fielded onto the NIF. Several issues affect the ability to acquire and lock onto Hartmann spots and allow the mirror to run through the calibration process required before operation. These include variations in average beam intensity, intra-beam intensity, and limited angular field of view. Since the interferometer used to characterize wavefront

control subsystems incorporates a Hartmann wavefront sensor, a video acquisition system, and a prototype NIF control system, it affords a unique opportunity to validate HS performance.

The prototype HS consists of a Cohu model 2122-2000 CCD camera and a 77-element lenslet array in which the focal length had been extended to 11.3 mm through the use of an index-adjusting fluid. NIF specifications call for a 77-element lenslet array with a 12-mm focal length using the same camera.

5.3.2.1 Intensity Variations of the Beam.

The first series of tests studied the power range over which the Hartmann wavefront sensor and video acquisition system could operate. The total beam power incident on the lenslet array was varied between 50 and 500 nW in the tests. The baseline estimate of power in the alignment beam at the HS is estimated to be limited to 205 nW. Three types of beams will be incident on the Hartmann wavefront sensor in the NIF: (1) the alignment beam, (2) the reference beam, and (3) the shot beam. The shot beam and the reference beam are expected to have plenty of power, which can be attenuated with ND filters.

The binary and gray centroiding thresholds, used by the video and acquisition systems to acquire and track Hartmann spots, were varied to find the range over which reliable spot acquisition and tracking occurred. Binary centroid calculations are used as an aid for initial spot acquisition, where all pixel values below the binary threshold are set to 0 and all pixel values above are set to 1. Gray-level centroid calculations are used during tracking of spots to improve noise immunity after initial acquisition. All values that are less than the gray threshold are set to 0, and all values above the threshold are set to their gray-level value.

Experiments showed that reliable acquisition and tracking occurred at power levels between 100 and 350 nW. At low-power levels, there was not enough light to acquire the spots. At high power levels, the light from Airy rings between the spots in the lenslet array focal plane (far field) started to reach the threshold intensity and to confuse the acquisition. Thus, the expected power level in the alignment beam is sufficient to acquire and track spots in the NIF, but the margin is less than a factor of two. (Note that during these experiments, the CCD utilized less than its entire dynamic range. This is not expected to have significantly affected the results of this test.)

5.3.2.2 Intra-beam Intensity Variations

A second series of tests were performed to determine how well the video acquisition system could acquire Hartmann spots when the power levels within the image varied, as expected in the NIF. Three types of beams will be incident on the Hartmann wavefront sensor in the NIF: (1) the alignment beam, (2) the reference beam, and (3) the shot beam. The alignment beam is expected to overfill the Hartmann array and will have little rolloff of the intensity near the edges of the

beam. It will, however, have a half-amplitude alignment spot at the center of the beam that covers the center lenslet. The reference source, which provides a beam to the Hartmann sensor that is free of the aberrations of beamline optics, is generated by mode-matching a single-mode fiber beam to the output lens of the transport spatial filter. This beam is expected to overfill the Hartmann array and will have a roll-off due to truncation of the Gaussian profile of about 50% near the corners of the beam. The beam during a shot will have little rolloff but will only half-fill the outside lenslets, even with perfect alignment. Therefore, in any of these cases, one would expect a factor-of-two difference in intensity between Hartmann spots within a given image. The worst case is when the intensity variations are abrupt, such as those within the alignment beam.

To test HS performance with such intra-beam intensity variations, an attenuation screen with 50% transmission was placed over a portion of the DM in the lab to attenuate the signal for nine Hartmann spots in one corner of the mirror. Experiments were then performed at various power levels while the binary and gray threshold values were varied to determine where reliable operation occurred. Only one power level, 250 nW, was found to result in reliable spot acquisition by the Hartmann wavefront sensor and video acquisition system. At this power level, the unattenuated Hartmann spots were operating with numerous pixels in saturation in each spot. The low-power limit was determined by the requirement to have enough power into the HS to acquire spots in the attenuated region. The high-power limit was determined by the need to keep the light between the Hartmann spots (overlapping Airy rings) in the unattenuated region below the acquisition threshold. It is expected that by employing the full dynamic range of the CCD, spot acquisitions will be successful over a wider range of operating power levels.

Thus, system performance is expected to be marginal with the baseline variations in intra-beam intensity. Possible approaches to mitigate this problem include the development of adaptive algorithms to set thresholds and to displace actuators to spread out the initial Hartmann spot positions in the regions where Airy rings between spots are overlapping to create bright spots that the image processor confuses with Hartmann spots. Reduction of the intensity difference of the beam-center alignment spot would also reduce this problem. These approaches are being investigated.

5.3.2.3 Hartmann Sensor Field of View

The third experimental series evaluated the impact on wavefront control system operation of a limited FOV in the optical system between the DM and the HS. This concern was raised by problems experienced in wavefront control systems fielded by other programs at LLNL, where light at high angles could not be read by the HS—particularly during control law calibrations when DM actuators are individually pushed and pulled. Beam field-stop apertures in the tower-to-output sensor package relay optics on the NIF will limit the FOV to about 200 μr . In addition, the diagnostic beamsplitter that sends the beam into the output sensor can

have up to three waves of focus and 0.6 waves of other aberrations in reflection at $1\text{ }\mu\text{m}$. As a result, the beam picked off from a perfectly corrected main beam will have this aberration present on it, which will further limit the range over which the mirror may be deformed before the beam is clipped by the limited aperture at the relay spatial filter.

To test HS performance with a NIF-like limited FOV, the FOV was limited on the interferometer in the wavefront control lab by placing apertures at the focus of the Keplerian beam-expanding telescope as shown in Figure 5-6. FOVs were limited to 211, 141, and $70\text{ }\mu\text{r}$, and tests were performed both with and without an aberrator placed into the 4-in. beam immediately after the splitter. This optic applied an astigmatic aberration of 4.3 waves P-V in transmission. A mirror calibration was performed with each of these FOVs by changing the voltage by $\pm 12\text{ V}$ centered at 90 V . The actual movement of the mirror depends on the individual actuator and its location, but a 200-V signal will displace the mirror surface from 2 to $7\text{ }\mu\text{m}$, which would result in between a 0.12- and $0.42\text{-}\mu\text{m}$ movement of the mirror surface during calibration. With no aberrator, the calibration routine operated reliably for FOVs of 211 and $141\text{ }\mu\text{r}$ but failed at the $70\text{-}\mu\text{r}$ FOV when a corner spot was clipped at the limited far-field aperture. With the aberrator installed, the calibration routine operated reliably for fields of view of 211 and $141\text{ }\mu\text{r}$, but was unable to even acquire Hartmann spots at the $70\text{ }\mu\text{r}$ FOV due to the clipping of the beam at the aperture. Therefore, FOV limits on the NIF are not expected to be a problem, but the margin is small.




Figure 5-6. Experimental setup for limited FOV trials.

5.4 Wavefront Controller

5.4.1 Wavefront Controller Design Concept and Validation Status

The NIF wavefront control computer system consists of two types of front-end processors (FEPs), as shown in Figure 5-7. The Image Processor FEP is based on a SPARC AXi processor, and the two wavefront control FEPs employ Motorola MVME 2306 processors. The Image Processor views four HS images and provides spot positions to the controllers via a dedicated ethernet link. The Controllers process the HS positions for each beam and calculate the appropriate actuator commands for each DM. The actuators are driven by Analog I/O carrier cards and by LLNL-designed amplifiers. The system employs modular hardware and object-oriented software. The NIF facility expected lifetime is 30 years, and software and hardware changes are expected over that time. By using a modular hardware and software architecture, system maintainability is improved significantly.

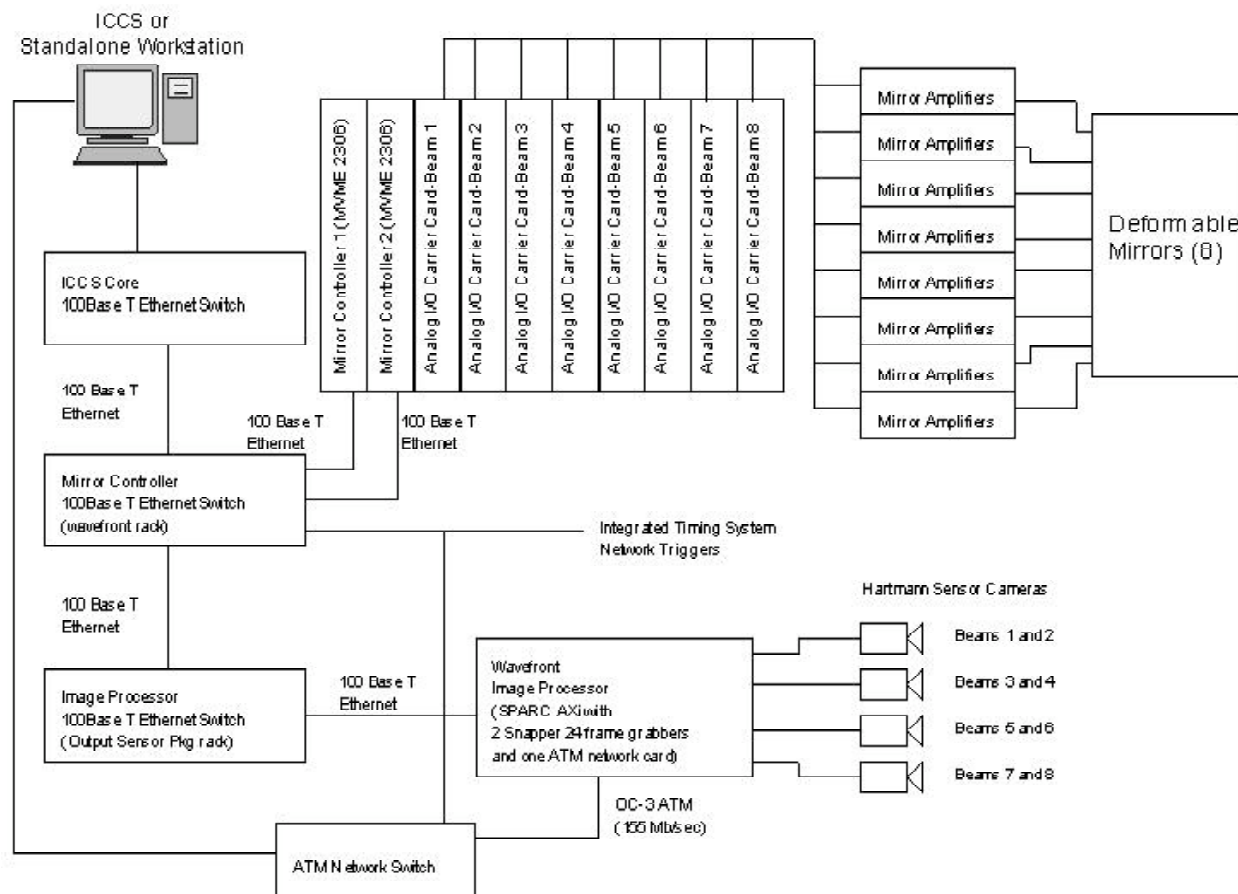


Figure 5-7. The NIF Wavefront Control System block diagram.

To achieve the 0.5 Hz system bandwidth requirement, the wavefront sensor is read at a 10-Hz rate. A planned read-out upgrade to 30 Hz should allow the goal of 1-Hz system bandwidth to be achieved. The sensor output is standard RS-170 video, which is read by an Active Imaging Snapper24 frame grabber. The digitized image is fed into a SPARCengine AXI computer that calculates centroids for all 77 lenslet spots and then calculates their offsets from the reference positions. This information is sent via a dedicated ethernet line to the Motorola MVME 2306 controller that calculates the required DM actuator displacements. Each image processor services four Hartmann sensors and eight beams. Each control computer services four DMs.

5.4.2 Wavefront Controller Risk Mitigation

A wavefront control system was built and tested as part of the LLNL DM interferometer at LLNL in FY98. The success of this one-beam controller resolved many risk areas of the controller. At the beginning of FY99, the remaining risk issues included operating bandwidth and cost. These issues were addressed by quantifying the per-mirror processor timing in the wavefront controller FEPs and by simplifying the DM actuator drive amplifiers.

5.4.2.1 Timing Analysis

During the past year, the software execution times of both wavefront control system FEPs were examined to verify whether the CPUs have sufficient speed to meet the design requirement (10-Hz closed-loop operation) and the design goal (30-Hz operation).

The critical operating mode for real-time operation is closed-loop wavefront control. All other system modes have open and flexible execution time deadlines (at least within the time for reasonable responsiveness to operator commands). To meet the 10-Hz requirement, all closed-loop control code (that supports and computes a single control-law loop iteration) on an FEP must execute within 100 milliseconds; to meet the 30-Hz requirement, all code must execute within 33.3 milliseconds.

Measurements indicate that the mirror controllers easily meet the requirement. On the MVME 2604 platform (an older NIF-baseline model that has since been superseded by the faster MVME 2306), the mirror controller closed-loop code executed within 2–3 milliseconds for a single beam (indicating 8–12 milliseconds for 4 beams/processor), well within both the 10- and 30-Hz requirements.

The system computational bottleneck is the Hartmann processor. Measurements made with a variation of the Beamlet C-based code (modified to run a single beam on the NIF hardware platform) indicated that total execution time for a single closed-loop iteration (the blob tracking code in this case) is about 10 milliseconds. This converts to a closed-loop iteration time for 8 beams/platform of about 80 milliseconds, within the 10-Hz requirement.

For the Hartmann processor to meet the 30-Hz requirement, closed-loop code for 8 beams would need to execute within 33 milliseconds. For a single beam, this reduces to execution within approximately 4 milliseconds. This goal has not been demonstrated yet. We expect that faster processors will be available before NIF Hartmann processing FEPs are procured, which will allow the goal to be met. Another option is to double the number of image processors; this would cost roughly an additional \$250K for the NIF.

The one other issue discovered during Video FEP development (during the Penlight milestone) was that acquisition and delayering of the red, green, and blue layers of an RGB image is a very CPU-intensive operation, taking around 24–30 milliseconds. For comparison, acquisition of a single-layer grayscale image takes 2–3 milliseconds. Since the delayering operation caused a large penalty in CPU time, it was decided to modify the Hartmann processor architecture to use four Snapper24 frame grabbers instead of two. This results in each camera having its own dedicated frame grabber card, and thus no delayering of multi-layer images is required. In effect, the delayering is done in hardware, by virtue of the platform architecture.

5.4.2.2 Actuator Drive Amplifier Simplification

The original NIF baseline PMN driver circuit was copied from that of a working wavefront control system developed for use on AVLIS, Beamlet, Petawatt, etc. It was designed to drive up to 16 A at 400 V, at up to several kHz. The NIF requirements are 0.25 A at 150 V, at up to 30 Hz, which is about a 200-times power reduction. To achieve its drive power requirement, the original design utilized a parallel-FET push-pull output stage that sat on top of a very expensive custom-made heat sink. Since the NIF drive power requirement is much less, elimination of the heat sink and elimination of the parallel FETs was possible. The original driver used standard through-hole components and could drive eight actuators per board, which meant that five boards would be required to operate one DM (39 actuators are required per mirror). With the lower power requirement, surface-mount components could be used instead of through-hole components, which resulted in a very space-efficient design. Each new design board can operate 20 actuators, which means that only two boards are needed to operate one DM. The net result is that only 400 boards will be needed for the the NIF instead of 1000.

A VME chassis is required to house the driver boards and processors. Each VME rack operates eight DMs. The old design that needed 40 amplifier boards for 8 DMs required a three-level VME chassis. With the new design, only 16 boards are needed. This means that a much cheaper two-level VME chassis can be used. The new amplifier design also includes a test connector that is expected to help reduce testing costs. The expected total savings resulting from these design changes is expected to be about \$650K.

References

1. R. A. Zacharias, et. al., "The National Ignition Facility (NIF) Wavefront Control System," *Proceedings of SPIE, Third Annual International on Conference Solid State Lasers for Application to Inertial Confinement Fusion*, SPIE **3492**, pp. 678–692, Monterey, California, 1998.
2. R. A. Zacharias, et. al., "Wavefront Control of High-Power Laser Beams in the National Ignition Facility," *Proceedings of SPIE, Advanced High-Power Lasers and Applications*, to be published, Osaka Japan, November 4, 1999.
3. R. Sacks, et al., "Application of Adaptive Optics for Controlling the NIF Laser Performance and Spot Size," *Proceedings of SPIE, Third Annual International on Conference Solid State Lasers for Application to Inertial Confinement Fusion*, SPIE **3492**, pp. 344–354, Monterey, California, 1998.
4. W. H. Williams, et al., "The NIF's Basic Focal Spot for ICF Temporally-Shaped Pulses," *Supplement to the Proceedings of SPIE, Third Annual International on Conference Solid State Lasers for Application to Inertial Confinement Fusion*, SPIE **3492**, pp. 55–60, Monterey, California, 1998.
5. W. H. Williams, et al., "NIF's Basic Focal Spot for Temporally Flat Pulses," *Supplement to the Proceedings of SPIE, Third Annual International on Conference Solid State Lasers for Application to Inertial Confinement Fusion*, SPIE **3492**, pp. 22–38, Monterey, California, 1998.
6. B. M. Van Wonterghem, et al., "Beamlet Pulse Generation and Wavefront Control System," *Inertial Confinement Fusion Quarterly Report*, **5**, Number 1, Lawrence Livermore National Laboratory, Oct.–Dec. 1994.
7. R. Hartley, et al., "Wavefront Correction for Static and dynamic aberrations to within 1 second of the system shot in the NIF Beamlet demonstration facility," *Proceedings of SPIE, Second Annual International on Conference Solid State Lasers for Application to Inertial Confinement Fusion*, October 22–25, 1996 Paris, France.
8. R. A. Zacharias, *National Ignition Facility Wavefront Control Development and Manufacturing*, 1999, NIF-0054450, Lawrence Livermore National Laboratory, September 2000.
9. J. F. Anzellotti, D. J. Smith, J. J. Sczupak, and Z. R. Chrzan, "Stress and environmental shift characteristics on HfO₂/SiO₂ multilayer coatings," in *Laser-Induced Damage in Optical materials: 1996*, H. E. Bennett, A. H. Guenther, M. R. Kozlowski, B. E. Newman, M. J. Soileas, Eds., Proc. SPIE **2966**, 258–264 (1997).

10. A. C. Erlandson and L. C. Truelson, *Fluence Measurements in the Polymer Edge Cladding Damage Tester*, LLNL internal memorandum ADG 86-16/2415R, March 7, 1986.
11. J. T. Salmon, et al., "Active and Adaptive Optical Systems," *Proceedings of the SPIE*, **1452**, pp.459–467, Bellingham, WA, 1991.
12. J. S. Toeppen, E. S. Bliss, T. W. Long, S. T. Salmon, "A Video Hartmann Wavefront Diagnostic that Incorporates a Monolithic Microlens Array," *Proceedings of SPIE, International Symposium on Optical Applied Science and Engineering*, July 21, 1991, San Diego, California.
13. E. S. Bliss, et al., "Design Process for NIF Laser Alignment and Beam Diagnostics," *Proceedings of SPIE, Third Annual International on Conference Solid State Lasers for Application to Inertial Confinement Fusion*, SPIE **3492**, pp. 285–292, Monterey, California, 1998.

6.0 NIF OPTICS MANUFACTURING	6-1
6.1 Properties of and Manufacturing Methods for NIF Laser Glasses.....	6-1
NIF Laser Glass Properties and Specifications	6-3
Spectroscopic and Optical Properties.....	6-3
Gain, Stored Energy, and Extraction Efficiency.....	6-3
Nonradiative Energy Losses	6-5
Nonlinear Refractive Index.....	6-7
Optical Quality.....	6-8
Properties of the NIF and LMJ Laser Glasses	6-11
Description of Glass Melting.....	6-12
Old Technology: One-at-a-Time, Discontinuous Melting and Forming.....	6-13
Advanced Technology: Continuous Laser Glass Melting and Forming	6-15
Postprocessing.....	6-17
Manufacturing Schedule.....	6-17
Summary	6-18
Acknowledgments	6-18
Notes and References.....	6-18
6.2 KDP Crystals.....	6-21
6.2.1 Growing KDP and DKDP Crystals for the NIF Laser.....	6-21
Introduction	6-21
Crystal Growth.....	6-23
KDP/DKDP Production for the NIF.....	6-35
References.....	6-36
6.3 Developing Optics Finishing Technologies for the National Ignition Facility.....	6-38
The NIF Optics Finishing Development Plan.....	6-38
Shaping, Grinding, and Polishing.....	6-39
Deterministic Figuring and Precision Polishing.....	6-43
Compensating for Material Inhomogeneities.....	6-48
Finishing for a High 3ω Laser-Damage Threshold.....	6-51
Notes and References.....	6-57
6.4 NIF Optic Processing	6-59
6.4.1 Introduction.....	6-59
6.4.2 Optics Processing Facility	6-59
Cleaning and Coating Equipment	6-61
Cleaning Systems.....	6-64
Coating Systems.....	6-66
Sol Coating Improvements.....	6-68
Protection of Sol Coatings during Service.....	6-71
References	6-76
6.5 Engineering High-Damage-Threshold NIF Polarizers and Mirrors	6-77
6.5.1 Thin-Film Deposition Technology.....	6-77
6.5.2 Electron-Beam Deposition.....	6-77
6.5.3 Ion-Beam Deposition Methods	6-79
6.5.4 Laser Damage.....	6-80
6.5.5 Delaminates.....	6-81
6.5.6 Flat-Bottom Pits.....	6-83
6.5.7 Nodular-Ejection Pits.....	6-85

6.5.8 Laser Conditioning to Minimize Damage	6-87
6.5.9 Plasma Scalds	6-88
6.5.10 Current Technology Status	6-89
6.5.11 Summary	6-90
Notes and References	6-91
6.6 Diffractive Optics for the NIF.....	6-92
6.6.1 The Beam Sampling Grating.....	6-94
6.6.2 The Color Separation Grating.....	6-95
6.6.3 The Effects of Sol-Gel Antireflective Coatings	6-96
6.6.4 Combining Diffractive Structures on One Surface.....	6-101
6.6.5 Summary.....	6-102
Notes and References.....	6-103

6.0 NIF Optics Manufacturing

6.1 Properties of and Manufacturing Methods for NIF Laser Glasses

J. H. Campbell, T. Suratwala, C. Thorsness, P. Ehrmann, W. Steele, M. Riley

The laser systems for the U.S. Department of Energy (DOE) National Ignition Facility (NIF) at Lawrence Livermore National Laboratory (LLNL) and for the French Commissariat à l'Energie Atomique (CEA) Laser Megajoule (LMJ) consist of 192 and 240 laser beamlines, respectively.¹⁻⁴ Each beamline contains either 16 (NIF) or 18 (LMJ) large, Nd-doped laser glass slabs, and each finished laser glass slab is about $81 \times 46 \times 4.1 \text{ cm}^3$. A total of 3072 and 4380 slabs will be installed on the NIF and LMJ, respectively (Table 6.1-1). In addition, each facility plans to purchase approximately 10% more slabs to be used as construction and operation spares. Thus, nearly 8150 laser glass slabs will be needed for the two laser systems, representing a volume of about 125 m^3 (330 metric tons) of finished, high-optical-quality glass. The quantity of raw glass that must be melted is greater than 330 metric tons to account for various processing losses.

Table 6.1-1. Quantity of laser glass required for the LMJ and NIF projects.

Variable	NIF	LMJ
Number of beamlines	192	240
Number of slabs/beamline	16	18
Spares	10%	10%
Total number of slabs (including spares)	3380	4752
Finished slab dimensions (cm^3)	$81 \times 46 \times 4.1$	$81 \times 46 \times 4.1$
Volume per slab (L)	15.3	15.3
Total volume (m^3)	52	73
Mass (metric tons)	136	190

Both the NIF and LMJ use a compact laser amplifier design called the multisegment amplifier (MSA).⁵⁻⁷ Amplifiers consist of stacked 4×1 arrays of laser glass slabs inside a flashlamp-pumped cavity (Figure 6.1-1). By using square apertures (i.e., square beams), it is possible to tightly pack the individual laser glass amplifiers into a compact matrix and greatly reduce the size and cost of the system.

This design requires that the laser glass be manufactured in rectangular slabs. Although the laser aperture is square, the laser slabs are rectangular because they are mounted at Brewster's angle to the propagation direction of the beam. Mounting the glass at Brewster's angle minimizes the Fresnel reflection losses at the surfaces of the slabs. In addition, mounting at an angle increases the coupling efficiency of the flashlamp light with the slabs. Erlandson et al.⁵⁻⁷ have recently described in detail the design and operating characteristics of flashlamp-pumped MSAs. The measured small-signal gain coefficient is typically about 0.05/cm, and the stored energy density is about 0.25 J/cm³ for phosphate laser glass doped at about 4×10^{20} /cm³ and pumped at a lamp explosion fraction of 0.20.



Figure 6.1-1. Assembly drawing of two 4×2 amplifier units to be used on the NIF. The $81 \times 46 \times 4.1$ -cm³ laser slabs are held in a frame that orients them at Brewster's angle relative to the propagation direction of the beam.

A prototype laser closely resembling one of the NIF beamlines has recently been built and tested.⁸ This laser, called Beamlet, uses 11 amplifiers in the main cavity and 5 in the booster section, the same configuration as that for the NIF. A series of large phosphate glass slabs ($767 \times 428 \times 44$ mm³) having a doping of 3.5×10^{20} /cm³ were produced for this laser.⁹ Although slightly smaller than required for the NIF and LMJ, these prototype glass slabs were made to nearly identical specifications as those required for the NIF and LMJ. Therefore, to a great extent, the quality and size of the laser glass pieces needed for the NIF and LMJ have been demonstrated. What remains is to develop the manufacturing capability for producing a large number of these high-quality glass slabs at a high rate and at significantly lower cost. This article briefly describes the advanced melting methods that are being developed to produce the glass for the NIF and LMJ. Although details of the production processes are proprietary, we highlight the work involved in changing from the one-at-a-time,

discontinuous production process used in the past to the continuous melting process of the future. Our cost goal is to manufacture the laser glass for about \$1000/L (\$350/kg), or roughly a factor of 3 lower than the cost associated with the current, one-at-a-time production methods. Given the size of the laser glass order, this level of cost reduction represents a total cost savings of about \$200 million to \$300 million compared to the price that could be achieved with current manufacturing methods.

NIF Laser Glass Properties and Specifications

To meet the performance requirements for NIF target applications, the laser glass must have the following important characteristics:

- High net gain.
- Efficient energy storage at high energy density.
- Efficient extraction of stored energy.
- Resistance to laser-induced damage.
- High optical homogeneity to allow light to propagate with negligible wavefront aberration.

Of these five properties, the first three relate to spectroscopic and optical properties of the glass, and the last two to glass quality.

Spectroscopic and Optical Properties

It can be argued that the most important laser glass spectroscopic and optical properties are the emission cross section, fluorescence lifetime, and nonlinear index of refraction. The emission cross section and fluorescence lifetime largely control the gain, energy storage, and extraction efficiency of the laser glass. The nonlinear index strongly impacts beam quality, particularly noise growth at the high operating intensities required for the NIF.

Compared to repetition-rated lasers that fire at frequencies of several hertz, the NIF will be essentially a single-shot device with a few hours between shots. Therefore, the physical properties of NIF glass tend to be of less importance for the end-use application. Nevertheless, the physical properties of glass can be very important to the successful manufacturing and handling of NIF laser glasses.

Gain, Stored Energy, and Extraction Efficiency

The amplification of a laser pulse passing through a laser gain medium (e.g., laser glass) of length z can be modeled using the well-known Frantz–Nodvik equation.¹⁰ This equation relates the output fluence F_{out} to the input fluence F_{in} , given the small-signal gain G_0 and saturation fluence F_{sat} of the laser medium, as follows:

$$F_{\text{out}} = F_{\text{sat}} \ln\{1 + G_0[\exp(F_{\text{in}}/F_{\text{sat}}) - 1]\} . \quad (1)$$

The small-signal gain of the laser glass is described by

$$G_0 = \exp\{z[(\sigma N^*) - \alpha]\} , \quad (2)$$

where σ is the emission cross section (cm^2), N^* is the Nd-ion inversion density ($1/\text{cm}^3$), and α is the transmission loss coefficient (cm^{-1}). Achieving high net gain requires a high emission cross section coupled with a large population inversion density. In addition, transmission losses due to absorption by impurities or scattering from defects in the glass (or on the polished surfaces) must be kept low relative to the gain.

The term σN^* is the gain coefficient g_L , which is related to the stored energy density ($E_s = h\nu N^*$) and the saturation fluence ($F_{\text{sat}} = h\nu_l / \sigma$) of the material:

$$g_L = N^* \sigma = (h\nu_l N^*) (\sigma / h\nu_l) = E_s / F_{\text{sat}} , \quad (3)$$

where h is Planck's constant ($\text{J} \cdot \text{s}$), and ν_l is the laser frequency (Hz). Experimentally, the gain coefficient g_L is measured directly by monitoring the signal gain of a probe beam passing through the glass as it is being pumped in a laser amplifier.^{11,12} For example, the large, flashlamp-pumped, glass laser amplifiers used in fusion energy applications typically operate with a gain coefficient of about 0.05 cm^{-1} . The emission cross section is determined independently from spectroscopic measurements. The phosphate glasses used in fusion applications nominally have an emission cross section of about 3.5 to $4.0 \times 10^{20} \text{ cm}^2$, corresponding to a saturation fluence of about 5 J/cm^2 . The stored energy in the laser glass is computed from Eq. 3 to be about 0.25 J/cm^3 .

Energy extraction from the glass is most efficient at high laser fluences, particularly fluences in excess of twice the saturation fluence. Having a glass with a high emission cross section is desirable because efficient extraction can be achieved at lower fluence, thereby reducing the chance of laser-induced damage to the optics.

Unfortunately, not all of the stored energy in the glass can be extracted, even at very high fluences.¹³ This limitation arises from the inhomogeneous broadening of the Nd emission. The cross section calculated from measurements of gain saturation using the Frantz–Nodvik equation give higher emission cross sections (σ_{gs}) than those determined spectroscopically (σ_{em}). The Nd-site inhomogeneities in the glass lead to hole burning, causing the glass to saturate at a lower fluence, corresponding to a higher effective cross section. The extraction efficiency then can be defined by the ratio of the two cross sections:

$$\eta_{\text{ext}} = \sigma_{\text{em}} / \sigma_{\text{gs}} . \quad (4)$$

Phosphate laser glasses tend to saturate much more homogeneously than do silicates; therefore, the efficiency with which the energy is extracted is much higher.¹³ This is one of the reasons for the wide use of phosphate glasses in ICF applications. To be more specific, a large fraction of the cost of large fusion research lasers is in the cost of the amplifiers. Therefore, simultaneously achieving high stored energy and efficient extraction is essential to minimize the number and cost of the amplifiers.

A high Nd^{3+} gain cross section is often desirable for ICF laser applications; however, it can also lead to some undesirable effects. For example, for large-aperture laser systems, a high gain coefficient can lead to low energy-storage efficiency and large spatial variations in the gain distribution across the aperture due to amplified spontaneous emission. The general rule of thumb is that the product of the gain coefficient and longest dimension of the laser glass piece should not exceed a value of about 4.2. In addition, the Nd doping and fluorescence lifetime must be optimized to achieve maximum energy-storage efficiency. The tradeoff among high gain, good spatial gain uniformity, and efficient energy storage is quite complex and is usually optimized with the use of sophisticated laser-design computer codes.¹⁴

Nonradiative Energy Losses

In most practical applications, the quantum yield is never 100% because of nonradiative relaxation mechanisms that significantly shorten the lifetime. The nonradiative losses are affected by intrinsic properties of the laser glass as well as the care with which the glass is manufactured.^{15,16} Figure 6.1-2 shows a schematic view of the most important nonradiative mechanisms.

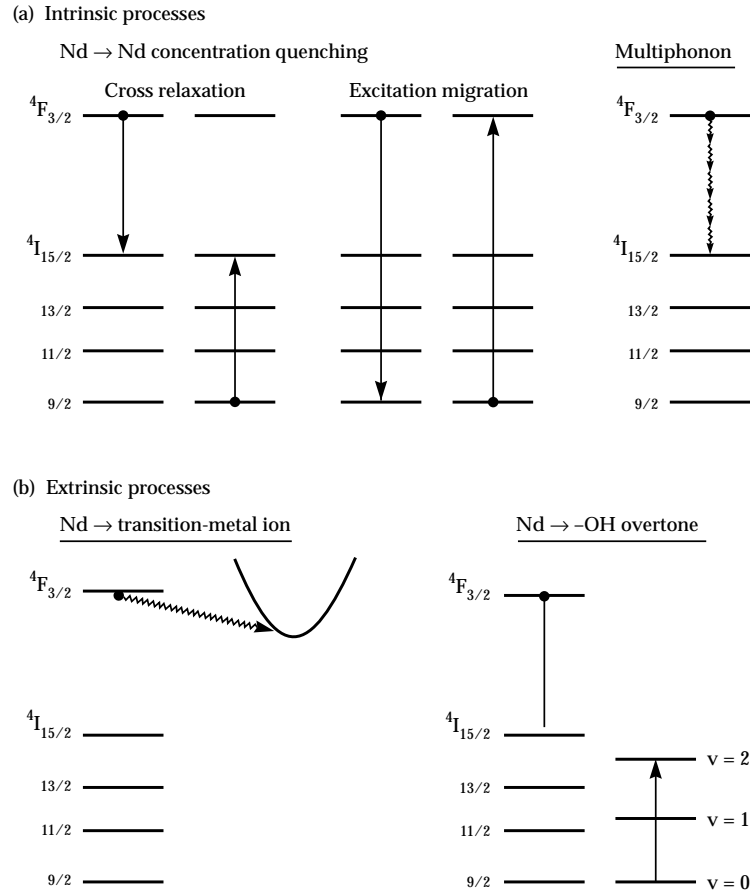


Figure 6.1-2. Nonradiative Nd³⁺ relaxation processes in laser glasses can be divided into (a) intrinsic processes that depend on glass composition and structure and (b) extrinsic processes that depend on impurities in the laser glass. Extrinsic losses depend strongly on processing methods.

The two most important intrinsic processes are multiphonon relaxation and concentration quenching. Concentration quenching refers to the radiation exchange between a pair of Nd ions. Nd concentration quenching results from the contributions of two relaxation mechanisms. One is cross relaxation in which the two ions share the energy, and the second is migration of the excitation energy from one ion to the next (the so-called “hopping” mechanism). In theory, the rate of relaxation due to concentration quenching varies as $1/r^6$, where r is the inter-ion distance, which is equivalent to the square of the Nd³⁺ concentration. In reality, the effect of the Nd concentration on relaxation rate is highly glass dependent and varies from nearly linear to quadratic in phosphates and from quadratic to cubic in silicates.^{15–18} Stokowski,¹⁵ and more recently Payne et al.,¹⁸ have proposed an empirical relation to characterize concentration quenching in metaphosphate laser glasses:

$$\tau_{em} = \tau_0 / [1 + (N/Q)^2], \quad (5)$$

where τ_0 is the zero concentration lifetime (μs), N is the Nd^{3+} ion concentration (ions/ cm^3), and Q (ions/ cm^3) is an empirically determined quantity for a given glass. Q is physically equivalent to the Nd concentration needed to reduce the lifetime τ_{em} to one-half its zero concentration limit.

The rate of multiphonon relaxation depends on how closely the $^4\text{F}_{3/2}$ -to- $^4\text{I}_{15/2}$ energy transition (5500 cm^{-1}) matches the maximum vibration energy of the glass matrix.^{15,16,19} For phosphate glasses, the maximum vibration energy is about 1170 to 1200 cm^{-1} (References 20 and 21), thus suggesting the absence of any significant multiphonon effects. Careful measurements by Caird et al.,¹⁶ on two commercial phosphate glasses (LG-750 and LG-760) support this observation. The measured multiphonon relaxation rate is about 200 s^{-1} .

It is well known that impurities that enter the glass during manufacturing (particularly hydroxyl groups and transition-metal ions) can cause significant nonradiative relaxation. Some useful correlations have recently been reported for estimating the impact of various impurities on the nonradiative relaxation rate in phosphate glasses.²² We hasten to add, however, that the problems associated with impurities have been largely eliminated in modern laser glass manufacturing by the proper choice of raw material and melting equipment.

Nonlinear Refractive Index

The energy extraction efficiency for ICF laser systems is also limited by nonlinear propagation effects, particularly at very short pulse lengths. The refractive index increases with the laser intensity:

$$n = n_0 + \gamma I , \quad (6)$$

where γ is the nonlinear refractive index coefficient (m^2/W), and I is the laser intensity (W/m^2). The intensity-dependent index can cause amplitude ripples (noise) that occur at certain spatial frequencies to grow exponentially:

$$I = I_0 \exp(B) , \quad (7)$$

where the B factor (or breakup integral) is the cumulative nonlinear phase retardation over the optical path length:

$$B = 2\pi/\lambda \int \gamma I \, dz . \quad (8)$$

Experience has shown that B needs to be less than about 2 rad to avoid unacceptable noise ripple growth.^{23,24} Such growth can cause optical damage and/or

degrade the beam focus. Therefore, laser glasses with low nonlinear indices are required for ICF applications.

Direct measurement of γ is difficult, so empirical correlations have been developed. The expression developed by Boling et al.²⁵ many years ago accurately predicts γ from the refractive index n_d and the Abbe number v of the glass:

$$\gamma = K(n_d - 1)(n_d^2 + 2)^2 / \{n_d v [1.52 + (n_d^2 + 2)(n_d + 1)v/6 n_d]^{1/2}\}, \quad (9)$$

where $K = 2.8 \times 10^{-10} \text{ m}^2/\text{W}$ is an empirically determined constant. The nonlinear refractive index n_2 (in esu) is related to γ by $n_2 = \gamma(nc/40\pi)$, where c is the speed of light.

Optical Quality

Some of the most critical specifications of laser glass relate to its optical quality, and the relevant characteristics are strongly dependent on the processing conditions. In particular, three main characteristics of the glass impact optical quality: optical homogeneity, inclusions, and bubbles.

Optical homogeneity refers to the refractive index variation in the optical material. For laser glass, the value is typically less than 2 ppm (i.e., $\Delta n < \pm 2 \times 10^{-6}$). Homogeneity is generally specified in terms of a maximum amount of allowed aberration due to sphere, astigmatism, and a smaller amount of higher-order terms (see Table 6.1-2). For the NIF and LMJ, we intend to keep the same specification. However the final, finished (i.e., polished) laser glass will be specified using a more sophisticated procedure designed to monitor aberrations at specific spatial frequencies that are known to seed nonlinear growth of intensity noise in the laser beam.

The homogeneity of the laser glass is critical to maintain wavefront uniformity of the laser beam. Recall that there are a total of 16 laser slabs in the laser beamline, and during multiple passes through the cavity and booster amplifier sections, the beam passes through the equivalent of 54 laser slabs. Therefore, even small optical inhomogeneities can lead to significant wavefront aberration in the output of the beam, potentially causing significant degradation in both frequency conversion and focusability of the beam.

Table 6.1-2. Key technical specifications for the NIF prefinished laser glass slabs. Prefinished slabs will be clad, ground, and polished into finished slabs by the finishing vendor.

Parameter	NIF specification
Nd doping	$4.2 \times 10^{20} \pm 0.1 \text{ Nd}^{3+}/\text{cm}^3$
Homogeneity (expressed as wavefront error at 632 nm, normal incidence)	
–Sphere	$\leq 0.425 \lambda$
–Astigmatism	$\leq 0.220 \lambda$
–Higher-order aberrations	$\leq 0.142 \lambda$
–Rms phase gradient (filtered 33 mm)	$\leq 0.0083 \lambda/\text{cm}$
Fluorescence lifetime (measured on $5 \times 5 \times 0.5 \text{ cm}^3$ sample)	$\geq 320 \mu\text{s}$
Absorption coefficients	
–At 1053 nm	$\leq 0.0019 \text{ cm}^{-1}$
–At 400 nm (due to Pt^{II+})	$\leq 0.25 \text{ cm}^{-1}$
–At 3333 nm (due to hydroxyl groups)	$\leq 2 \text{ cm}^{-1}$
Bubbles	
–Maximum number (per 100 cm^2 area)	Total cross section $< 0.15 \text{ mm}^2$
–Maximum diameter	$\leq 1000 \mu\text{m}$
Birefringence	$\leq 5 \text{ nm}/\text{cm}$
Pt inclusions	
–Maximum number for any one slab	≤ 5 in clear aperture
–Maximum size after laser irradiation (5 shots per site)	$\leq 750 \mu\text{m}$

Inclusions from ceramic refractory materials, unmelted raw materials, platinum (Pt) metal, crystallites, or impurities can cause optical damage in the glass when exposed to high laser fluences. The most common inclusion source is metallic Pt inclusions from the Pt liners used in the melting system. Improved processing conditions have led to a dramatic reduction in Pt inclusions in recent years such that the average inclusion density is less than 0.1 per liter of glass, or less than an average of 1 to 2 per glass slab.^{26,27} Inclusions in laser glass typically damage at about 2 to 5 J/cm² at the NIF and LMJ pulse lengths.²⁸ Although very small to begin with, inclusion damage can grow with successive laser shots to several millimeters or even centimeters in size, making the laser glass unusable. Large damage spots in laser glass can seed damage in other optics in the laser chain. In general, if the inclusions are small, they can be tolerated as long as the optical damage they produce does not exceed about 750 μm in size. This limit is the basis for the specification given in Table 2. Currently, we scan each piece of laser glass with a

high-fluence laser beam and measure the size of any damage site after a specified number of shots at fluences between 7 to 14 J/cm² (8 ns).²⁹⁻³¹ If the Pt-damage size remains below the specified size limit given in Table 6.1-2, then it is acceptable.

New processing techniques for eliminating Pt inclusions rely on the intrinsic property of many phosphate glasses to dissolve Pt metal under oxidizing conditions. Izumitani et al.³² have reported the use of gaseous POCl₃ as an oxidizing additive, and Campbell and coworkers at Hoya Corporation (Fremont, Calif.) and Schott Glass Technologies, Inc. (Duryea, Pa.)²⁶⁻²⁸ have reported the effects of O₂, Cl₂, CCl₄, and N₂/O₂ mixtures. Model predictions of Pt-inclusion dissolution rates agree well with results from glass-melting tests.³³

In general, glasses with the largest platinum solubility have the lowest inclusion density. The effects of glass composition on Pt solubility have been reported to follow the trend: phosphate > silica-phosphate >> fluorophosphate > silicate.³² Thus, phosphates have much lower inclusion densities than do silicates. The solubility trend is based on measurements using LHG-5 and LHG-8 (phosphates), HAP-3 (silica-phosphate), LHG-10 (fluorophosphate), and LSG-91H (silicate).

If inclusions are eliminated from the bulk laser glass, the damage threshold is limited only by the quality of the surface finish.³⁴ This is shown in Figure 6.1-3, where surface damage threshold is plotted versus pulse length for various laser glasses as well as fused silica. The pulse-length dependence of the surface damage threshold for finely polished samples can be accurately represented by the empirical expression

$$D_s \text{ (J/cm}^2\text{)} = 22t_p^{0.4}, \quad (10)$$

where t_p is the pulse length (ns). The surface damage threshold approximately follows the $t^{1/2}$ relation predicted by a thermal-diffusion heat-transport model. For comparison, the damage threshold reported by Gonzales and Milam³⁵ for Pt inclusions is also shown in Figure 6.1-3. It follows the approximate pulse-length scaling relation

$$D_{Pt} \text{ (J/cm}^2\text{)} = 2.5t_p^{0.3}. \quad (11)$$

Thus, the presence of Pt inclusions reduces the operating limit of the laser glass by nearly tenfold.

The specification for laser glass bubbles is based on two requirements. The first is the need to minimize the amount of light loss due to obscurations caused by bubbles. The second is the need to keep the size below a certain value that may induce nonlinear growth of intensity noise. The obscuration loss is not to exceed 0.01% of the beam area per slab and therefore sets the total number of bubbles of a given size allowed in any given slab. The maximum-size bubble allowed is currently 1000 μm. The diffracted light from bubbles that exceed 1000 μm can, at high intensities, imprint a holographic diffraction pattern in the next optic that, in turn,

can bring that portion of the beam to focus at another downstream optic and potentially damage it. Because of the regular spacing of many optics in the laser chain, such a nonlinear imaging effect could lead to propagation of laser damage throughout the beamline. In general, bubbles have not been a significant problem for laser glass. For example, only one of the Beamlet slabs had bubbles, and they were so few in number and so small as to be insignificant.⁹ Similar results were observed for the Nova and Phebus laser glass disks.




Figure 6.1-3. Measured surface-damage threshold for various polished laser glasses and fused silica vs pulse length at 1053 nm. The bulk-damage thresholds also shown arise from Pt inclusions, as reported by Gonzales and Milam.³⁵

Properties of the NIF and LMJ Laser Glasses

Laser glasses are specially formulated to give the desired laser, optical, thermal-mechanical, and physical-chemical properties needed for a specific laser application. Some properties are strongly affected by the processing conditions as discussed above; however, most are controlled by the base glass composition.

We have chosen two glasses for use on the NIF and LMJ that meet the gain, energy-storage, extraction-efficiency, and damage-resistance requirements: LHG-8 (Hoya Corporation) and LG-770 (Schott Glass Technologies, Inc.). LHG-8 is the same glass as that used on the Nova and Phebus lasers. However, LG-770 is a new formulation developed to replace the LG-750 glass that was used on Nova, Phebus, and Beamlet. Table 6.1-3 summarizes the key properties of these laser glasses. Figure 6.1-4 shows the neodymium absorption and emission spectra for the two glasses, confirming that the glasses are essentially indistinguishable. In fact, for nearly all properties, the two glasses closely match, thereby making it possible for us to

randomly interchange the two glass types within the NIF amplifiers without affecting laser performance.

Table 6.1-3. Properties of LHG-8 (Hoya) and LG-770 (Schott) laser glasses.

Glass properties	Symbol	LHG-8	LG-770
Optical			
Refractive index at 587.3 nm	n_d	1.5296	1.5067
Refractive index at 1053 nm	n_1	1.5201	1.4991
Nonlinear refractive index			
n_2 (10^{-13} esu)	n_2	1.12	1.01
γ (10^{-20} m ² /W)	γ	3.08	2.78
Abbe number	v	66.5	68.4
Temp. coefficient refract. index (10^{-6} /K)	dn/dT	-5.3	-4.7
Temp. coefficient optical path (10^{-6} /K)	δ	0.6	1.2
Laser			
Emission cross section (10^{-20} cm ²)	σ_{em}	3.6	3.9
Saturation fluence (J/cm ²)	F_{sat}	5.2	4.8
Radiative lifetime (zero Nd) (μ s)	τ_0	365	372
Judd-Ofelt radiative lifetime (μ s)	τ_r	351	349
Judd-Ofelt parameter (10^{-20} cm ²)	Ω_2	4.4	4.3
Judd-Ofelt parameter (10^{-20} cm ²)	Ω_4	5.1	5.0
Judd-Ofelt parameter (10^{-20} cm ²)	Ω_6	5.6	5.6
Emission bandwidth (nm)	$\Delta\lambda_{eff}$	26.5	25.4
Concentration quenching factor (cm ⁻³)	Q	8.4	8.8
Fluorescence peak (nm)	λ_L	1053	1053
Thermal			
Thermal conductivity (W/mK)	k	0.58	0.57
Thermal diffusivity (10^{-7} m ² /s)	α	2.7	2.9
Specific heat (J/Kg)	C_p	0.75	0.77
Coefficient thermal expansion (10^{-7} /K)	α_e	127	134
Glass transition temperature (C)	T_g	485	461
Mechanical			
Density (g/cm ³)	ρ	2.83	2.59
Poisson's ratio	μ	0.26	0.25
Fracture toughness (MPa • m ^{0.5})	K_{IC}	0.51	0.43
Hardness (GPa)	H	3.43	3.58
Young's modulus (GPa)	E	50.1	47.3

Description of Glass Melting

Glasses made for the present, large ICF laser systems (e.g., Nova at LLNL, Phebus at CEA, Beamlet at LLNL, Gekko at Osaka, and OMEGA at the University of Rochester) were manufactured using a one-at-a-time, discontinuous melting process. We briefly describe this older manufacturing method and then discuss the new, advanced processes that have been developed for the NIF and LMJ.





Figure 6.1-4. The (a) absorption and (b) emission spectra for LG-770 and LHG-8 laser glass are virtually indistinguishable.

Old Technology: One-at-a-Time, Discontinuous Melting and Forming

The first step of the discontinuous process is premelting (Figure 6.1-5), which is designed to melt and mix the raw starting materials. A bubbling gas is often added to

remove unwanted volatile products, particularly water, and if necessary, to adjust the melt redox state. The premelt is carried out in a relatively inert refractory crucible. The walls of the refractory vessel corrode over time, eventually requiring the vessel to be replaced. Glass from the premelter generally contains bubbles, striae, and possibly some small particles of unmelted starting material.



To complete one full melt cycle requires about 1 to 2 days

Figure 6.1-5. Schematic of the current, discontinuous, two-step process used to melt and form laser glass slabs.

Product glass from the premelt stage is next processed in a physically separate unit called the remelter. The remelter consists of a platinum-lined vessel that also has provisions for stirring and gas bubbling. The main purposes of the remelter are to dissolve any Pt inclusions, remove any bubbles, and homogenize the glass to provide the striae-free, high-optical-quality glass necessary for laser applications. This process involves several stages.³⁶ During the first stage of the remelt cycle, the redox state of the glass is adjusted to enhance Pt-particle dissolution. Next, a refining or “fining” process is conducted at high temperatures, where the viscosity of the glass is low, allowing bubbles to rise to the surface. The third stage is a stirring process that is generally conducted at temperatures lower than those for either the melting or refining stages. Continuous stirring thoroughly distributes all components within the glass melt, eliminating striae and thus ensuring uniformity of the refractive index over the entire casting. Finally, the melt is cooled to a temperature such that the viscosity of the glass is proper for casting into a mold of

the appropriate size and shape. After casting, the glass undergoes a coarse annealing step, is inspected for inclusions and striae, and then is fine-annealed to remove residual thermal stresses due to the forming process. Although it only takes one or two days to melt and cast each slab, it can take several months to complete the subsequent annealing and inspection.

Advanced Technology: Continuous Laser Glass Melting and Forming

Advanced laser glass melting processes have been developed separately by Schott Glass Technologies, Inc., and Hoya Corporation under work funded jointly by LLNL and the Centre d' Etudes de Limeil-Valenton. The two glass companies have chosen different development approaches. Schott has chosen to design and develop a full-scale melting system that will then become the production melter. The Schott approach allowed for one development run to verify equipment design and the melting and forming process. Such verification was completed in November and December 1997. In contrast, Hoya chose to carry out development using a subscale, continuous melter. Because of the smaller size and lower operating costs of their equipment, Hoya was able to carry out several melting and forming campaigns, which were completed in March 1998. Both vendors have completed their development efforts, and as of February 1999, they are in the midst of the first full-scale run that we term the "pilot." This work will be followed by several years of production, approximately 2 to 3 years for the NIF and about 3 to 4 years for LMJ.

Many details of the manufacturing process are highly proprietary to each company. Therefore, we give only a generic description of the melting, forming, and coarse annealing process. Nevertheless, this description should provide an idea of the progress in laser glass manufacturing technology that has occurred as a result of the NIF and LMJ projects. The laser glass melting systems developed by Schott and Hoya are arguably the most advanced optical glass melting systems in the world.

As shown in Figure 6.1-6, a continuous optical glass melting system is generally divided into several interconnected zones. Each zone consists of one or more vessels designed to carry out a specific aspect of the process. For laser glass continuous melters, Figure 6.1-6 shows the six main processing zones,³⁷ which are raw material batching, melting, conditioning, refining, homogenizing, and continuous strip forming. The six regions are interconnected, allowing for a continuous flow of glass from one zone to the next.

It is desirable that the raw materials be batched together and then thoroughly mixed in a dry atmosphere. The batch is then delivered continuously to the melter with precautions to avoid water uptake by hygroscopic raw materials. Batch powder that enters the melter dissolves in the molten glass and undergoes large-scale mixing. Off-gas handling equipment collects any gas emissions from the melter (or other vessels) and treats the effluent to meet environmental regulations.

Glass continuously flows from the melter into the conditioning unit, where the redox state of the melt is adjusted to enhance dissolution of Pt inclusions. If required, steps may also be taken to remove any excess “water” (i.e., hydroxyl groups) in the glass. Glass from the conditioning unit then flows to the refiner section where the temperature is generally elevated to reduce the glass viscosity and thereby increase the bubble rise velocity to promote bubble removal.




Figure 6.1-6. Schematic of the continuous laser-glass melting systems to be used to manufacture NIF and LMJ laser glass. The six process zones are discussed in the text.

Glass from the refiner enters the homogenizing section, where Pt stirrers thoroughly mix the glass to achieve the part-per-million index homogeneity required for ICF laser applications. Just as in the discontinuous process, the temperature of the homogenizing section is reduced to adjust the glass viscosity to give the desired flow characteristics needed to form a wide, thick, homogeneous strip of glass. The width and thickness of the glass strip produced during the forming operation are greater than those of any optical glass ever produced prior to NIF- and LMJ-driven glass development.

The glass manufacturers employ highly proprietary technology to “form” (i.e., cast) the glass into a homogeneous, continuous strip free of sharp index variations (striae). Once successfully formed, the cast strip moves by conveyor belt through a long (25- to 35-m), coarse annealing oven (Figure 6.1-7) where the temperature is ramped down at a rate to avoid generating unacceptable thermal stresses in the glass. Finally, the cast strip is cut into pieces that are individually processed to give the desired laser slab blank.

Both manufacturers will use advanced processing conditions designed to minimize the formation of Pt inclusions in laser glass. Prior to 1986, Pt inclusion

damage represented the major source of damage in laser glass used for high-peak-power applications. However, new processing methods effectively reduce the Pt inclusion concentration by more than 1000-fold, to fewer than an average of 1 to 2 per laser glass slab (i.e., less than 0.1 per liter).




Figure 6.1-7. Schematic of the system used to coarse-anneal the as-cast, continuous laser glass strip.

Postprocessing

Once the laser glass has been melted and formed into plates, several other process steps must be completed before the glass can be shipped to the final finishing vendor. Specifically, the laser glass needs to undergo prefabrication to a size suitable for inspection for striae and Pt inclusions. Next, the glass is slowly annealed to remove any residual strain, a process that can take many days. Finally, the glass blank is fabricated to the final dimensions, inspected for homogeneity, and prepared for shipping.

Manufacturing Schedule

Laser glass manufacturing for the NIF and LMJ is divided into two main phases: pilot and production. Pilot refers to the first production runs. Results from the pilot run will be used to establish yield and costs. In addition, glass from the pilot runs will be used by the finishing vendors to demonstrate the advanced laser glass finishing and polishing methods to be used in final production.

The production phase immediately follows the pilot phase. The first stage of laser glass production will be primarily for the NIF facility because NIF construction will occur earlier than that for the LMJ. The NIF production will take place over approximately three years at a rate of about 1000 slabs per year. NIF production will

be followed by LMJ production, which will last three to four years. During the third year of production, the NIF and LMJ may overlap somewhat, requiring a short-term increase in the annual production rate.

Summary

The NIF and LMJ laser systems require about 3380 and 4752 Nd-doped laser glass slabs, respectively. Continuous laser glass melting and forming will be used for the first time to manufacture these slabs. Two vendors have been chosen to produce the glass: Hoya Corporation and Schott Glass Technologies, Inc. The laser glass melting systems that each of these two vendors have designed, built, and tested are arguably the most advanced in the world. The cost goal is to manufacture the laser glass for about \$1000/L (\$350/kg), or roughly a factor of 3 lower than the cost associated with current, one-at-a-time production methods. Production of the laser glass began on a pilot scale in the fall of 1998.

Acknowledgments

The authors gratefully acknowledge J. Hayden, A. Thorne, H. Pankratz, J. Cimino, and their colleagues at Schott Glass Technologies, Inc., and K. Takeuchi, M. Smolley, K. Suzuki, J. Storms, and their colleagues at Hoya Corporation for their fine efforts in developing the advanced laser glass processes needed to produce the laser glass for the NIF and LMJ. Our close collaboration with Ms. Gaëlle Ficini-Dorn and her colleagues at the French CEA is gratefully acknowledged.

Notes and References

1. J. Paisner and J. R. Murray, "Overview of the National Ignition Facility Project," *Proc. SPIE Conf. Solid-State Laser Apps. to ICF* (Monterey, CA, 1998).
2. M. Andre, "LMJ and LIL," *Proc. SPIE Conf. Solid-State Laser Apps. to ICF* (Monterey, CA, 1998).
3. *National Ignition Facility Conceptual Design Report*, vol. 2 and 3, Lawrence Livermore National Laboratory, Livermore, CA, UCRL-PROP-117093 (May 1994).
4. M. Andre, *Chocs, Revue Scientifique et Technique de la Direction de Applications Militaires* **11**, 82–85 (July 1994).
5. A. C. Erlandson, M. D. Rotter, D. N. Frank, and R. W. McCracken, *Inertial Confinement Fusion Quarterly Report* **5**(1), 18–28, Lawrence Livermore National Laboratory, Livermore, CA, UCRL-LR-105821-95-1 (1994).
6. A. Erlandson et al., "Flashlamp-Pumped Nd:Glass Amplifiers for the National Ignition Facility," *Proc. 13th Embedded Topical Meeting on the Technology of Fusion Energy* (American Nuclear Society, 1998).

7. L. Zapata et al., "Gain and Wavefront Measurements Performed on the NIF/LMJ Prototype Amplifiers," *Proc. SPIE Conf. Solid-State Laser Apps. to ICF* (Monterey, CA, 1998).
8. B. M. Van Wonterghem et al., *Applied Optics* **36**, 4932 (1997).
9. J. H. Campbell et al., *Inertial Confinement Fusion Quarterly Report* 5(1), 29–41, Lawrence Livermore National Laboratory, Livermore, CA, UCRL-LR-105821-95-1 (1994).
10. L. M. Frantz and J. S. Nodvik, *J. Appl. Phys.* **34**, 2346–2349 (1963).
11. H. T. Powell, A. C. Erlandson, K. S. Jancaitis, and J. E. Murray, *High Power Solid State Lasers and Applications, SPIE Proc.* **1277**, 103–120, Bellingham, WA (1990).
12. A. C. Erlandson, K. S. Jancaitis, R. W. McCracken, and M. D. Rotter, *Inertial Confinement Fusion Quarterly Report* 2(3), Lawrence Livermore National Laboratory, Livermore, CA, UCRL-LR-105821-94-3 (1994).
13. W. E. Martin and D. Milam, *IEEE J. Quantum Electron.* **OE18**, 1155–1163 (1982).
14. *National Ignition Facility Conceptual Design Report*, vols. 2 and 3, Lawrence Livermore National Laboratory, Livermore, CA, UCRL-PROP-117093 (May 1994).
15. S. E. Stokowski, *Laser Glass: An Engineered Material*, Lawrence Livermore National Laboratory, Livermore, CA, UCRL-96331 (1987).
16. J. A. Caird, A. J. Ramponi, and P. R. Staver, *J. Opt. Soc. Am. B* **8**, 1391 (1991).
17. S. A. Payne, C. D. Marshall, A. Bayramian, G. D. Wilke, and J. S. Hayden, *Appl. Phys. B* **61**, 257–266 (1995).
18. S. A. Payne et al., *Ceramic Transactions: Solid State Optical Materials* **28**, 253–260, American Ceramic Society Press (1992).
19. C. B. Layne, W. H. Lowdermilk, and M. J. Weber, *Phys. Rev. B* **16**, 10 (1977).
20. H. Toratani, *Properties of Laser Glasses*, Ph.D. Thesis, Kyoto University, Japan, (1989).
21. H. Toratani, I. Izumitani, and H. Kuroda, *J. Non-Cryst. Solids* **52**, 303–313 (1982).
22. J. H. Campbell and T. I. Suratwala, *J. Non-Cryst. Solids* **263**, 318–341 (2000).
23. J. T. Hunt and D. R. Speck, *Optical Engineering* **28**, 461–468 (1989).
24. J. T. Hunt, K. R. Manes, and P. A. Renard, *Appl. Optics* **32**, 5973–5982 (1993).
25. N. L. Boling, A. J. Glass, and A. Owyong, *IEEE J. Quantum Electron.* **QE14**, 601 (1978).
26. J. H. Campbell, E. P. Wallerstein, J. S. Hayden, D. L. Sapak, and A. J. Marker, *Glotech. Ber. Glass Sci. Technol.* **68**(1), 11–21 (1995).
27. J. H. Campbell, E. P. Wallerstein, H. Toratani, H. Meissner, and T. Izumitani, *Glotech. Ber. Glass Sci. Technol.* **68**(2), 1–11 (1995).
28. J. H. Campbell et al., *Elimination of Platinum Inclusions in Phosphate Laser Glasses*, Lawrence Livermore National Laboratory, Livermore, CA, UCRL-53932 (1989).
29. C. L. Weinzapfel et al., "Large-Scale Damage Testing in a Production Environment," *Laser Induced Damage in Optical Materials: 1987* (NIST Special Publication 756, National Institute of Standards and Technology, 1987), 112–122.
30. J. H. Campbell, J. F. Kimmons, and S. Schwartz, "Platinum Particle Detection in Phosphate Laser Glass," *Analysis of the Composition and Structure of Glass and Glass Ceramics* (Springer-Verlag, Heidelberg, 1999), Chapter 6.6.

31. S. Schwartz et al., "Vendor-Based Laser Damage Metrology Equipment Supporting the National Ignition Facility," *Proc. SPIE Conf. Solid-State Laser Apps. to ICF* (Monterey, CA, 1998).
32. T. Izumitani, M. Matsukawa, and H. Miyade, "Solubility of Pt in Nd Phosphate Laser Glass," *Laser Induced Damage in Optical Materials: 1987* (NIST Special Publication 756, National Institute of Standards and Technology, 1987), pp. 29–34.
33. J. H. Campbell, *Glass Sci. Technol.* **68**, 91–101 (1995).
34. J. H. Campbell et al., "Damage Resistant Optics for a Mega-Joule Solid-State Laser," *Laser Induced Damage in Optical Materials: 1990, SPIE 1441*, 444–456 (1990).
35. R. Gonzales and D. Milam, *Laser Induced Damage in Optical Materials: 1985*, NBS Special Publication 745 (1988), pp. 128–137.
36. A. J. Marker, "Optical Glass Technology," *Geometrical Optics, SPIE Proc.* **531**, 2–10 (1985).
37. J. H. Campbell, T. I. Suratwala, C. B. Thorsness, J. S. Hayden, A. J. Thorne, J. Cimino, A. J. Marker III, K. Takeuchi, M. Smolleg, G. F. Fincini-Dorn, *J. Non-Cryst. Solids* **263**, 342–357 (2000).

6.2 KDP Crystals

Alan Burnham and Ruth Hawley-Fedder

The cost and physics requirements of the NIF have established two important roles for potassium dihydrogen phosphate (KDP) crystals: beam polarization rotation in the main laser cavity to enable a multipass architecture and frequency conversion in the final optics assembly to convert the 1053-nm fundamental frequency of a neodymium glass laser to 351 nm, which couples better to the targets. The polarization rotation uses a KDP switch crystal in a Pockels cell, and the frequency conversion uses a KDP doubler and DKDP tripler crystals. The location of these crystals in the laser is shown in Figure 6.2-1.

Manufacturing of these crystals for the NIF faces three challenges: growing a crystal boule large enough and with sufficient quality to make the crystal plates; fabricating the plates with the required flatness, smoothness, orientation, and surface damage resistance; and coating the plates with an anti-reflective layer that does not damage the surface of the crystal. The first challenge is discussed in Section 6.2.1—boules of both KDP and DKDP have been grown with the required size and quality. The UV optical damage aspect of crystal quality is discussed in further detail in Section 4.2.4. The second challenge, fabricating the plates, is discussed in Sections 4.1.1–4.1.3. In addition, possible contamination from finishing as a source of optical surface damage is discussed in Section 4.2.4. Although crystal plates meeting NIF specifications had been made in small numbers for Beamlet, the new fabrication and finishing production equipment nearing completion will be able to make the slightly larger NIF crystals more efficiently. While the general manufacturing issues for antireflective coatings are discussed in Section 6.4, the specific challenges relating to crystal coatings are discussed in Section 4.3. The major issue is protecting the crystal surface from moisture that can condense in the porosity of a sol-gel coating. Major progress has been made in using thermal annealing for switch crystals, barrier coatings for doubler crystals, and surface-modified sols for tripler crystals as means to minimize or eliminate this problem.

6.2.1 Growing KDP and DKDP Crystals for the NIF Laser

Ruth Hawley-Fedder, Harry Robey, Natalia Zaitseva, James De Yoreo, Alan Burnham

Introduction

The NIF requires KDP switch crystals, KDP doubler crystals, and DKDP tripler crystals, as shown in Figure 6.2-1. To accomplish their different functions, the orientation of the plates with respect to the crystal boule crystal axes are different. Consequently, different size boules are required for the three plate types, even for the same 41-cm size. The switch crystal is the easiest, since it is cut horizontally with respect to the base of the crystal. As a result, a 43-cm-square boule is large enough, including a 1-cm buffer for finishing purposes, and the plates stack efficiently up to or even into the pyramidal cap, depending on the size of the base. Only about 15 of these small boules

are needed. The KDP doubler crystal is the most challenging, as it is rotated in two axes with respect to the base. As a result, the minimum size base for a single doubler crystal is 51 cm, and a symmetric 55-cm-square by 55-cm-high boule will generate only 6 doublers, so nearly 35 of these boules would be required. A symmetric DKDP tripler must be a minimum of 55-cm high to be tall enough for a single tripler, but that minimum size will yield about 15 triplers.

Boules of both KDP and DKDP meeting NIF size and quality requirements have been grown by Cleveland Crystals, Inc., (CCI) by what is often called conventional growth. In this case, impurities in the growth solution poison growth of the vertical faces (prisms), thereby maintaining a cross-section approximately equal to that of the original seed crystal as the pyramid faces grow. Unfortunately, this pyramidal growth is very slow, and it takes about two years to grow a crystal to NIF size. To provide more programmatic flexibility and reduce costs in the long run, we have developed an alternative technology commonly called rapid growth. Through a combination of higher temperatures and higher supersaturation of the growth solution, a NIF-sized boule can be grown in 1–2 months from a small “point” seed. However, growing boules of adequate size is not sufficient. Care must be taken to prevent inclusions of growth solution and incorporation of atomically substituted impurities in the prism growth. Other issues important for meeting transmitted wavefront quality, absorption, and laser damage criteria must be addressed also.

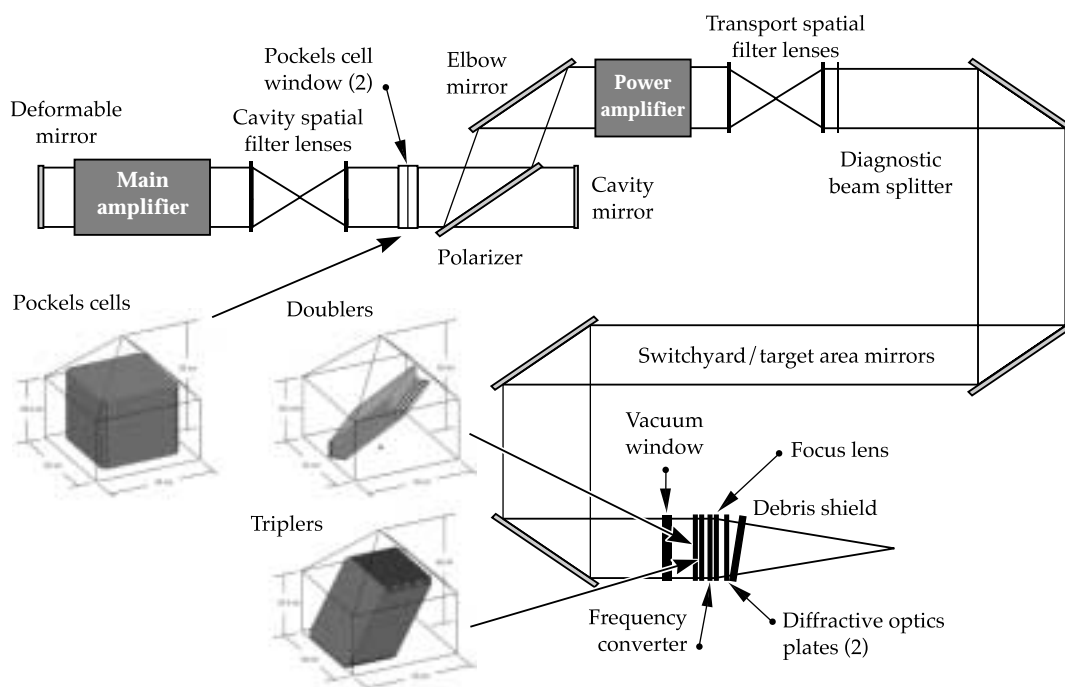


Figure 6.2-1. Schematic of the NIF laser showing the location of the KDP and DKDP components. Also shown is the orientation of the conversion crystals within the boules.

During the past year, we brought the rapid-growth technology to the stage needed to supply most of the KDP and some of DKDP for the NIF. This article reviews the technical hurdles that were overcome during this period and outlines some of the issues yet to be fully resolved. It also covers some of the efforts to transfer this technology to two potential vendors for NIF crystals, CCI and Inrad, as well as the development of crystal finishing technology.

Crystal Growth

For several years, the rapid-growth process has reproducibly grown crystals of high visual quality up to about 20 cm in linear dimensions in tanks containing less than 100 L of solution. These systems demonstrated many of the important conditions needed to grow crystals of NIF quality, such as solution purity and ways to achieve adequate mass transfer at the growing crystal surface. Over the past year, this understanding has been applied at the 1000-L tank scale in order to improve yields of NIF size boules to the point at which rapid growth can become an industrial production process. This section reviews the fundamentals of the growth process and the particular problems that affect the commercial viability of rapid growth for NIF-scale boules.

The fundamental property of nature enabling rapid crystal growth is that KDP can attain very large and stable supersaturations in solution. In other words, KDP will not spontaneously crystallize from solution when a solution is prepared at high temperature and then cooled so that the salt concentration is above its equilibrium solubility. This is because any proto-crystal formed by a statistical fluctuation has to reach a minimum size before growth is thermodynamically favorable. Measurements have shown that stable supersaturations from 35% at 65°C to 100% at 10°C can be attained if the solution is thoroughly preheated to eliminate any nucleation sites, and stability is not affected by impurities at the level of tens to hundreds of parts per million.¹ This supersaturation is much higher than the 3–20% required for growing crystal faces at rates of 10–20 mm/day.

When a seed crystal is introduced into this supersaturated solution, the crystal immediately grows at a rate that depends on a variety of chemical kinetic and mass transfer factors. Practical experience over many years has shown that the best way to start this process is to first partially dissolve an oriented seed crystal of about 1 cm³ in size above the saturation temperature, then decrease the temperature until the solution supersaturation reaches about 3%, at which time the seed “regenerates.”² Regeneration is a process in which a rectangular base and pyramid form over the partially dissolved, rounded seed crystal. A picture of a regenerated seed is shown in Figure 6.2-2. Each crystal face has numerous imperfections called dislocations. The crystal grows by adding atoms from solution to a set of atomic steps that emanate from these dislocations. This configuration is called a growth hillock, and a microscopic picture of a growth hillock is shown in Figure 6.2-3.³ As growth proceeds, the stronger growth hillocks crowd out the weaker ones, and favorable growth proceeds with one to four hillocks on each crystal face. Crystals continue to grow as long as the temperature is

decreased to maintain appropriate supersaturations. The growing crystal is rotated back and forth on a horizontal platform with a washing machine-like action to maintain good mass transfer.



Figure 6.2-2. Photograph of a recently regenerated seed. The cloudy regeneration layer is covered with about 1 cm of clear growth.




Figure 6.2-3. Atomic force microscopy picture of the atomic steps emanating from a dislocation on the crystal surface. The resulting macroscopic feature is called a growth hillock. The crystal grows by adding atoms from solution at the edge of the steps.

Even though the growth solution is stable with respect to homogenous nucleation, occasional heterogeneous formation of unwanted seeds at a variety of possible locations can cause the formation of unwanted crystals, which are generally first observed on the bottom of the tank. Unwanted seeds, whether from the original spurious source or from secondary crystals at the bottom of the tank, are transported throughout the

solution, and some will inevitably land on the product crystal, spoiling the surface and ruining the growth run. Eliminating these spurious crystals for the entire two months of a growth run has been one of the most important challenges for cost-effective production. During the three-year period from October 1995 through September 1998, only 12 growth runs out of 97 ended without spurious crystals. Since October 1998, 20 out of 36 growth runs have gone to completion without spurious crystals, and an additional 5 growth runs have been successfully completed in the presence of spurious crystals. Careful attention to details in the assembly of growth tanks and care in following procedures and equipment modifications to improve reliability have contributed to our success in minimizing the likelihood that spurious crystallization will result in the premature conclusion of a growth run. Although not completely eliminated, the occurrence of spurious crystals during a growth run is no longer a serious concern.

Growth difficulties related to solution impurity levels continue to be of importance. Impurities in solution are important for a variety of reasons: (1) they affect the growth rates of the prism faces, (2) they can enhance the formation of inclusions of growth solution that reduce optical quality, (3) they substitute into the atomic lattice in the prism sectors and cause inhomogeneities in the refractive index and loss of optical transmission by absorption, and (4) they may form particulates that affect the laser damage threshold of the material.

While the chemical structure of the pyramid face causes typical ionic impurities to be rejected from the growing crystal, the chemical structure of the prism face causes them to be selectively absorbed and incorporated into the crystal.⁴ A few examples of this selective absorption and rejection are shown in Table 6.2-4. Fe is the most important impurity in terms of transmittance, because FePO_4 is highly absorbing at 351 nm, resulting in a maximum acceptable Fe concentration of 100 ppb for salt from which DKDP triplers are grown.

Table 6.2-1. Concentrations of typical impurities in the raw material and in the pyramidal and prismatic sectors of KDP crystals.

Impurity	Raw material*	Pyramid*	Prism*
B	1000	ND	ND
Na	86000	ND	ND
Al	900	200	4400
Si	12000	<100	390
Ca	3600	ND	ND
Cr	2000	490	11000
Fe	5300	110	12000

*units of ng/g KDP

ND = not detected

Impurity incorporation into the prism face also affects the relationship between growth rate and supersaturation, as shown in Figure 6.2-4.⁵ At low supersaturations,

impurities effectively stop prism growth, and this condition is called “the dead zone.” Conventional growth occurs in this region, and only pure pyramid material is formed. Unfortunately, this corresponds to growth rates less than 1 mm/day, which results in growth times greater than one year for NIF-size boules. In the middle transition zone, relatively small changes in growth conditions can have a drastic effect on growth rate. As is typical for any stable, reliable industrial process, this region should be avoided. This leaves the high growth region as the most appropriate for development.




Figure 6.2-4. Effect of impurities on the step velocity and normal growth rate of KDP. The optimal growth rate is above the nonlinear region of step velocity and an upper bound determined by the uniformity of mass transfer.

While the effect of impurities on growth rates is understood quantitatively over some range of conditions, it is not understood quantitatively for mixtures of impurities and over the wider temperature range used to grow crystals. Rather than attempt to quantify this parameter space in detail, our approach was instead to reduce impurities to the lowest practical level, which is also important for meeting other specifications, and to use the qualitative principle of maintaining the highest possible growth rate to minimize the effect of impurities on growth instability. Once the necessary purity of the starting salt was attained (<0.5 ppm impurities), the contribution of impurities from the Pyrex growth tanks was explored. Though usually inert, there is a finite rate of tank dissolution in the hot KDP solutions used for rapid growth.⁶ Using measured Pyrex dissolution rates and the uptake coefficients for various impurities in the prism sector, we have successfully modeled the buildup and eventual consumption of Al and Fe in the growth solution, as shown for Al in Figure 6.2-5. Elements such as B and Si are not absorbed in the crystal and continue to build up in the growth solution, while others such as Ca are roughly constant during each growth run but increase after each resaturation of the growth solution. Polycarbonate tanks are being evaluated in calendar 2000 to eliminate this problem. Another important recent advance for

maintaining crystal quality is the successful implementation of constant filtration to remove particulates that come from the moving equipment or from precipitation.⁷




Figure 6.2-5. Comparison of measured and calculated aluminum concentration in the growth solution as a function of time. The aluminum concentration initially rises as the rate of glass tank solution dominates when the crystal is small; it later falls because the rate of uptake in the prism face increases with crystal size and exceeds the glass dissolution rate, which is dropping with temperature.

At high growth rates, another problem becomes important. Variations in KDP concentration on the μm scale at the growing crystal steps can cause inclusions of growth solution. These inclusions can easily be large enough to cause more obscuration by scattering than can be tolerated in the laser. Figure 6.2-6 summarizes much historical data on the occurrence of pyramidal inclusions in 1000-L growth tanks through 1998, indicating that pyramidal inclusions are much more prevalent at high growth rates. A similar trend is believed to occur for prismatic inclusions. This presents the crystal grower with a dilemma—how to avoid the detrimental effects of impurity buildup and associated degradation of prismatic KDP that occurs at slow growth rates, while avoiding the potential formation of massive pyramidal inclusions at high growth rates.




Figure 6.2-6. Summary of 1000-L growth tank data showing the relationship between pyramidal inclusions and growth rates for slow and fast rotational regimes. For runs prior to late 1998, the rotational rate was typically 25 rpm, while the high rotation rate during the recent run represented by the triangles was at 50 rpm.

Microscopic investigations, hydrodynamic modeling, and theoretical modeling have been combined to provide a good mechanistic understanding of the formation of inclusions. Current theory suggests that there are two types of instability that contribute to the formation of inclusions. One type of instability involves the bunching of elementary growth steps into what are known as macrosteps.^{8,9} Macrosteps are the result of longitudinal instability on the crystal surface caused by depletion of KDP from solution near the growing surface. This depletion results in localized changes in individual step growth rates. Faster growing steps will overtake slow steps, resulting in the formation of macrosteps consisting of bunches of elementary steps. The formation of macrosteps can be controlled by changing the direction and speed of rotation of the growing crystal, as well as by changing the supersaturation. Macrosteps by themselves are usually not a problem, but when macrosteps from two sides of a growth hillock bend around and approach each other from the opposite direction, a deep valley can be formed.⁹ Inclusions also tend to form in the vicinity of this valley. Valley formation can be minimized by rapid rotation, thereby keeping the solution concentration in the

center of the crystal face closer in magnitude to that along the edges. Figure 6.2-7 shows the effect of increasing rotation on valley formation.

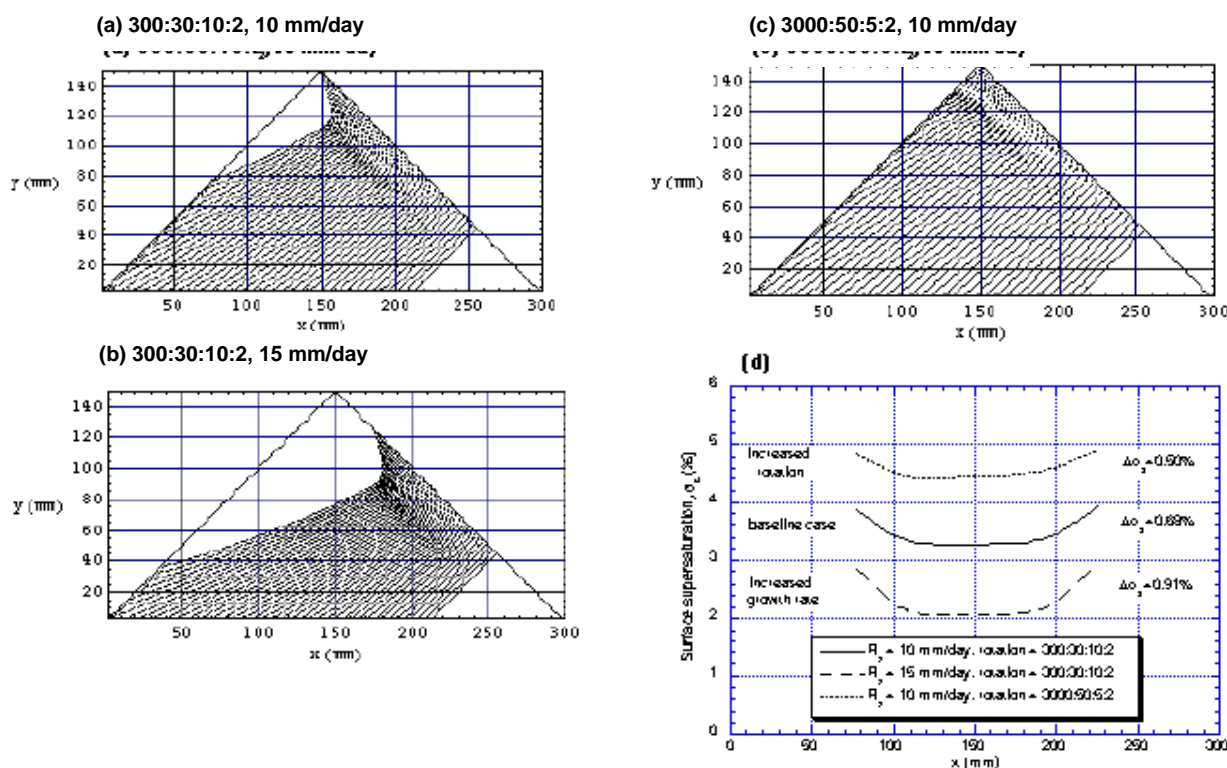


Figure 6.2-7. Effect of rotation and growth rate changes on valley formation and supersaturation across the pyramid face. A rotation schedule of 300:30:10:2 with a growth rate of 10 mm/day is taken as the baseline case (a), where 300 is the acceleration in rpm/m, 30 is constant rotation velocity in rpm, 10 is the time in seconds at constant velocity, and 2 is the pause time in seconds before changing rotation direction. For a constant growth rate of 10 mm/day, a change in rotation to 3000:50:5:2 (c) results in no step bending across the pyramid face and an increase in effective supersaturation across the crystal face (d). An increase in growth rate to 15 mm/day at the baseline rotation (b) results in an increase in step bending and a decrease in surface supersaturation.

A micrograph of a second type of unstable growth known as “fingering” is shown in Figure 6.2-8.⁸ Fingering is caused by lateral instability in which the supersaturation at the tips of the fingers is higher than the supersaturation experienced by the growing step in the cavities between the fingers. Although thermodynamics promote the filling of cavities between the fingers during slow growth, kinetics can cause the fingers to grow catastrophically, thereby surrounding and occluding growth solution. Solution moving in the opposite direction of the step advance is depleted in concentration as it moves past the tips of the fingers and towards the cavities, thereby causing a faster growth rate at the tip of the finger than at the base. Rapidly alternating the direction of flow helps prevent this catastrophic growth. Hydrodynamic modeling studies completed in early 1999 support the thesis that rapid rotation is necessary to provide more uniform feeding of the growing crystal surfaces. As the rotation is increased from

30 rpm to 50 rpm, the supersaturation experienced at the center of the growth face more closely approaches that of the bulk solution. In addition, the difference in supersaturation between the center of the crystal face and the edges of the crystal face decreases (see Figure 6.2-7). This is because the shear stress experienced in each region is significantly increased.




Figure 6.2-8. Fingering growth of macrosteps on the KDP surface.

Through a combination of these fundamental studies and growth experiments at various scales over the past year, we have shown that better mass transfer by increased acceleration and rotation rates can increase the inclusion-free growth rate by 20–40% over that shown in Figure 6.2-6. The best-quality NIF-size KDP boule grown up through late 1998 followed the solid-square growth trajectory also shown in Figure 6.2-6, substantially above that previously considered safe. An important aspect of achieving these higher rotation rates was the design and fabrication of streamlined aluminum growth platforms coated with a nonleaching Halar® coating. The combination of improvements in the growth hardware and increased understanding of the mechanisms of inclusion formation means that inclusions are now rarely seen in rapid-growth KDP and DKDP boules.

With the resolution of previous problems associated with spurious crystallization and inclusion formation, our efforts now have focused on defining the parameters that will allow us to reproducibly control the crystal geometry and aspect ratio. The geometry of the product boule is also a practical problem. The smaller size requirements associated with production of switch crystals means that boules that fail to meet doubler size can usually be used for switch crystal production. The requirements for doubler and tripler boules means that the aspect ratio, defined as the

ratio of boule height to average base width, is as important as the boule geometry in determining the final crystal yield from a boule. For symmetrical boules with dimensions of $55 \times 55 \times 55$ cm, the maximum yield is 7 doubler crystals (KDP) and 17 tripler crystals (DKDP). If the boule dimensions change to $50 \times 62 \times 62$ cm, corresponding to an aspect ratio of 0.8, the maximum yields are reduced to 6 doubler crystals and no tripler crystals. Because of the crystal geometry within the boule, DKDP must be grown to a height of 53.4 cm before any yield is realized. A minimum aspect ratio of 0.87 (53.4 cm tall, 61.4 cm wide) is the smallest symmetrical boule of DKDP that will still yield triplers with current rapid-growth equipment.

Regardless of boule geometry, the aspect ratio of a growing boule changes throughout the growth run. Immediately after regeneration, most boules display an aspect ratio close to 1. The aspect ratio then decreases steadily for the first 10–14 days of the growth run, or until the crystal base reaches between 10–12 cm. At this point, the aspect ratio begins to increase. This change corresponds to the point during growth when the prism faces enter the grooves of the growth platform. Figure 6.2-9 is a plot showing the change in aspect ratio for a series of NIF-sized boules grown at LLNL. When the growth rates of the various crystal faces are compared, it is seen that the observed increase in aspect ratio is due primarily to a reduction in prismatic growth rates.




Figure 6.2-9. Aspect ratio change as a function of growth run day.

No single factor has been identified as the primary cause of the decrease in prism growth rate. In fact, there are several competing factors that contribute to the observed change in the growth rate of the prism faces. At about the same time as the aspect ratio begins to increase, there is an observed change in the location of the growth hillocks on the prisms. As the prism faces enter the grooves, the dislocation structure of the faces changes and the growth hillocks move toward the center of the prisms. This movement of the hillocks must be due to the presence of a strong dislocation source at the grooves. A dislocation source strong enough to move growth hillocks as is observed would be expected to result in an increased growth rate for the prism faces. However, differences in supersaturation gradient across the prism face would be expected to result in decreased prism growth rates. Hydrodynamic modeling has shown that there is a supersaturation gradient across the crystal faces, which results in reduced feeding of the prism center as compared with the edges of the crystal. For a 30-cm crystal, the decrease in supersaturation ($\Delta\sigma$) from the edge to the center of the prism face has been calculated to be 0.11 to 0.22, and represents a 5–10% decrease in the salt available for growth. This means less salt is available for feeding of a central growth hillock. Solution impurities are also increasing until about day 25–30 of the growth run. Prismatic growth is especially sensitive to impurity concentration, and higher impurity levels that cause pinning of atomic steps on the surface are expected to lead to a decrease in growth rate on the prism faces.

Both aspect ratio and boule geometry have an effect on maximum crystal yield from a NIF-sized boule. Minor variations in crystal dislocation structure, combined with equally minor variations in hydrodynamic feeding of the growing crystal surfaces, usually result in asymmetrical crystals. The effects of asymmetry on crystal yield can be enormous. For doublers, the ideal asymmetry is one in which two adjacent prismatic faces grow taller than the other two prismatic faces. This type of asymmetry is illustrated in Figure 6.2-10. In this case, the observed asymmetry resulted in a 60% increase over the maximum yield for a symmetrical crystal of the same dimensions ($54 \times 56 \times 57$ cm, aspect ratio 1.036).




Figure 6.2-10. Rapid-growth KDP boule (photograph, left, and CAD drawing, right) showing desired asymmetry resulting in a yield of 11 doubler crystals. Maximum yield for a symmetrical boule of the same dimensions is 7 doublers.

When grown rapidly from a pure solution necessary to meet quality objectives, DKDP tends to grow with an aspect ratio lower than KDP. Given the limitations in the crystal base size imposed by the existing circular growth platform diameter of 90 cm, the height a symmetric-based crystal with an aspect ratio of 0.8 would reach is 51 cm high—too short for triplers. Improving the aspect ratio by allowing more solution impurities to retard prism growth is undesirable, especially for DKDP, which is challenged to meet a more difficult 351-nm laser damage threshold. As a result, we are exploring a variety of ways to improve both doubler and tripler yields by making more optimal shapes. A method for triplers is to grow the DKDP boule horizontally, which increases aspect ratio (now rotated) by eliminating one growth prism and adding a second pyramid. Three good-quality horizontal DKDP boules similar to that shown in shown in Figure 6.2-11 have been grown to NIF size.



Figure 6.2-11. Horizontal growth boule (DKDP) that will yield about 12 triplers.

Based upon data on growth dimensions collected during these runs, there appears to be a natural limit to the aspect ratios observed for horizontally grown DKDP. We monitored two aspect ratios associated with horizontal DKDP growth: the ratio of the height (prism z) to width (prism y) and the ratio of the width (prism y) to length (pyramid x). These dimensions are shown in Figure 6.2-19.




Figure 6.2-12. Horizontal boule dimensions monitored for aspect ratio determination.

The length/width ratio is important as it determines the maximum size boule that will fit on the platform. The height/width ratio is important for determining when a crystal is large enough to yield tripler crystals. For horizontal growth DKDP, the minimum height for yield is 22.5 cm, the minimum width is 43 cm, and the minimum length is 68 cm. Maximum dimensions are governed by the platform diameter (90 cm) and the space available between the platform bars for removal of the boule (63 cm). A plot of these aspect ratios is shown as Figure 6.2-12. The height/width ratio appears to reach a constant value between 0.34 and 0.40, where the minimum acceptable ratio is 0.357. The implication for horizontal growth is that boules will grow to fairly large base dimensions (length and width) before the boule height reaches the minimum necessary for yield. If the vertical prismatic growth rates can be increased to match the horizontal prismatic growth rates, then the height/width aspect ratio should improve. Some increase in the vertical prismatic growth rate should be possible by a slight modification of the tank configuration during growth. All data shown in Figure 6.2-13 was collected from growth runs that included a paddle above the growing boule that rotated in the same direction as the growing boule. Hydrodynamic modeling of the effect of a paddle shows that a paddle rotating with the crystal decreases the maximum shear stress experienced by the pyramid face by 15% when compared to the shear observed with no paddle. There is a corresponding decrease in the effective supersaturation seen at the center of the pyramid. Although these models were made for vertically grown boules, a similar effect would be expected for the upper growth surface of a horizontally oriented boule.




Figure 6.2-13. Observed aspect ratios for horizontal KD*P boules grown to date. Open symbols are length/width ratios; closed symbols are height/width ratios. The upper set of parallel lines represent the maximum and minimum length/width ratios and the lower set of parallel lines represent the optimum and minimum aspect ratios for height/width ratio.

KDP/DKDP Production for the NIF

Since 1998, LLNL has partnered with two commercial crystal growth vendors to transfer the rapid-growth technology necessary for the success of the NIF to the private sector. Cleveland Crystals, Inc. (CCI) located in Cleveland, OH and Inrad Inc. located in Northvale, NJ have each successfully grown NIF-size boules of KDP using the procedures developed and proven at LLNL. Pilot contracts were placed with both CCI and Inrad in late 1998 to enable them to demonstrate success using the rapid-growth techniques. Upon completion of the vendor pilot efforts, one-year contracts for continued production of KDP boules were placed with both CCI and Inrad.

In combination with growth efforts at LLNL, rapid growth has produced sufficient KDP doubler boules to produce almost half of the doublers and switch crystals needed for initial operation of the NIF. Development of production capabilities for DKDP has resulted in growth of four DKDP boules capable of producing enough tripler crystals for one quarter of the initial NIF needs. Our overall success rate is 66%, defined as the number of completed growth runs that produced boules of sufficient size and quality to meet NIF requirements (Figure 6.2-14).

This success rate is important when considering the operational lifetime of the NIF and the expected replacement rates for doubler and tripler crystals. We might need to produce sufficient DKDP to replace almost 50% of the tripler crystals each year of NIF operation, and enough KDP to replace as much as 10% of the doubler crystals. These replacement rates mean that during NIF operation, up to five rapid-growth tanks may be required to produce the necessary replacement crystals. This level of continuous production means that the success of our commercial vendors is critical to insure the continued operation of the NIF once construction is complete and all laser lines are installed.




Figure 6.2-14. Success of rapid growth for production of NIF-sized KDP and DKDP boules.

References

1. N. P. Zaitseva, J. J. De Yoreo, M. R. DeHaven, R. L. Vital, L. M. Carman, and H. R. Spears, *J. Crystal Growth* **180**, 255–262 (1997).
2. N. P. Zaitseva, I. L. Smolsky, and L. N. Rashkovich, *Krystallografiya* **36**, 198 (1991).
3. J. J. De Yoreo, T. A. Land, L. N. Rashkovich, T. A. Onischenko, J. D. Lee, O. V. Monovskii, and N. P. Zaitseva, *J. Crystal Growth* **182**, 442–460 (1997).
4. N. P. Zaitseva, L. Carman, I. Smolsky, R. Torres, and M. Yan, “The effect of impurities and supersaturation on the rapid growth of KDP crystals,” submitted to *J. Crystal Growth*.
5. L. N. Rashkovich and N. V Kronskey, *J. Crystal Growth* **182**, 434–441 (1997).
6. H. Robey, R. Floyd, R. Torres, and A. Burnham, “Impurity leaching rates of 1000 liter growth tanks,” Lawrence Livermore National Laboratory, Livermore, CA, UCRL-ID-133365.

7. N. Zaitseva, J. Atherton, R. Rozsa, L. Carman, I. Smolsky, M. Runkel, R. Ryan, and L. James, "Connection between continuous filtration and dislocation structure of KDP crystals," submitted to *J. Crystal Growth*.
8. H. F. Robey and S. Potapenko, "Ex-situ microscopic observation of the lateral instability of macrosteps on the surface of rapidly grown KH_2PO_4 crystals," submitted to *J. Crystal Growth*.
9. H. F. Robey, S. Potapenko, and K. Summerhays, "Bending of steps on rapidly grown KH_2PO_4 crystals due to an inhomogeneous surface supersaturation field," submitted to *J. Crystal Growth*.

6.3 Developing Optics Finishing Technologies for the National Ignition Facility

T. G. Parham, T. Baisden, C. Kiikka, C. Stolz, M. Kozlowski, D. M. Aikens

The National Ignition Facility (NIF), which will house a 192-beam laser system capable of generating two million joules of ultraviolet light, is being constructed at Lawrence Livermore National Laboratory (LLNL). To build this laser system, LLNL's industrial partners must produce more than 7500 meter-class optics. Their large precision-optics manufacturing facilities began producing the optics in 1998 and will finish their production by 2003.

In 1994, a series of cost studies made it clear that in order to fabricate (at the lowest possible cost) this unprecedented number of large precision optics in so little time, new technologies would have to be developed and new factories, based on them, would have to be constructed.¹ At that time, working with three suppliers of large optics, LLNL began three-year-long optics-finishing programs to develop these technologies. Although each development program centered on the specialties and often proprietary technologies already existing at the suppliers' facilities, many technologies for the low-cost manufacturing of these optics were common to two, or sometimes all three, of these programs.

Because many advances in these programs stem from the vendors' intellectual property and trade secrets, they cannot be completely described in a public forum. Nevertheless, many nonproprietary advances have been made and are now being used by the companies under contract: Eastman Kodak, SVG-Tinsley Division (formerly Tinsley Laboratories and referred to hereafter as Tinsley), and Zygo. The nonproprietary, optics-fabrication, developmental highlights for the manufacturing steps of shaping, grinding, polishing, figuring, and metrology are discussed² in a general sense in this article.

The NIF Optics Finishing Development Plan

The developmental effort in optics-fabrication technology was divided into three parts:

- A small internal LLNL effort (optics finishing R&D, wavefront analysis, and specifications).
- Three fabrication-development subcontracts for flats.
- Two fabrication-development subcontracts for aspheric lenses.

All vendor subcontracts were initiated by competitive solicitation. The three companies selected worked closely with LLNL to advance the technologies most critical to their own manufacturing processes. In most cases, successful subscale experiments were followed by full-scale experiments on production equipment. Table 6.3-1 lists the funded activities.

Overall, the program was highly successful; each vendor demonstrated a process that would satisfy all (or nearly all) of the NIF's performance requirements, at a fraction of the pre-NIF cost. The facilities construction has been completed, or is nearly complete, at all three vendors. They will yield about a ten-times-higher capacity for large precision optics than their previously existing facilities based on earlier technologies—and at prices that are roughly one-third the cost of equivalent Beamlet optics.

Table 6.3-1. Funded development activities.

Technology	LLNL	Vendors
Fixed-abrasive grinding	None	Hybrid-tool grinding; ELID dressing
Table lapping/loose-abrasive grinding	Double-sided lapping	High-speed, pellet-lap grinding
Polish out	High-pressure polishing; thermal figure control	High-speed polishing; variable-pressure polishing; synthetic-lap polishing
Deterministic figuring	None	Reduced-ripple, small-tool polishing; deterministic planetary polishing; ion-beam figuring; variable-pressure polishing
Metrology	PSD-based, wavefront specifications	Various
Other tools/processes	Cladding and cleaning specifications; automated cleaning	Optic handling; improved tooling; cladding technology
Laser-damage-performance improvements	Grinding-material selection; slurry; optimization; subsurface damage minimization; inspection-tool development	Slurry optimization; subsurface damage minimization

Shaping, Grinding, and Polishing

Although substantially more than half the fabrication cost of a precision optic accrues during the final figuring and metrology steps, the easiest way to reduce these costs is often before polishing—by improving the part's surface quality, smoothness, figure, and reducing the amount of subsurface damage. Consequently, about half of the development projects were in the areas of shaping, surface grinding, and the polish out of the optical surfaces to remove diffuse scatter before the first in-process optical test.

The first step in producing a NIF optical component is “shaping.” Many detailed variations on this process step exist, but the goals are all the same:

- Complete the outside dimensions to the specified length and width dimensions and tolerances.
- Produce an edge finish that can be precision cleaned (a rms surface microroughness of 0.2 μm).
- Generate the precision optical surfaces.

Optical surface-shape (or figure) control is critical. This is especially true during the early process steps, as most of the material is removed at this stage. Therefore, the ability to correct surface shape or orientation, relative to other optical surfaces, decreases as the optic is processed. Furthermore, if the optical surface shape is not maintained throughout processing, the depth of material removal at subsequent process steps will not be uniform. Residual subsurface damage could then reduce laser-damage thresholds in the areas of lesser material removal.

In the past, a two-part grinding process was used: a fixed-abrasive grinding step for coarsely shaping the optic to remove a large amount of material and a second step of loose abrasive grinding that is done in several stages, each with a smaller grit size and less material removed.

The NIF development vendors have adapted commercial computer numerical controlled (CNC) mills to shape the outside edges of the parts. One example is a hybrid rotary tool (Figure 6.3-1) developed by Tinsley for use on their CNC mill, which eliminates the need for loose abrasive grinding to the final surface roughness. It features two surface-generating areas on the tool, each with a different abrasive size. This allows removing large amounts of material in a single pass, with good surface figure control, while simultaneously producing a smooth, finished surface. Internal cooling channels are used to rapidly remove the heat being generated. Edging tests demonstrated removal rates consistent with optics production from saw-cut blanks using a single CNC edging machine. Their new machine now allows them to produce NIF lenses and windows for a lower overall cost, since lower-cost blanks with saw-cut surfaces can now be purchased, instead of blanks with finely ground surfaces.



Figure 6.3-1. The hybrid grinding/shaping tool developed by Tinsley to improve the efficiency of initial blank shaping.

Zygo is using a recent advance in fixed abrasive grinding, a technique called Electrolytic In-process Dressing (ELID), to produce smoother optical surfaces with fewer process steps. Typically, flat or spherical optical surfaces are generated using a

coarse, bound abrasive wheel—followed by successively finer, loose-abrasive grinds. Published reports on the ELID technique indicated that near-specular finishes could be obtained from a fine, bound, abrasive over small apertures.³ This work indicated the existence of a technology that might be applicable to large optics to eliminate at least one loose-abrasive grinding step, or possibly more.

The major restriction of conventional fixed-abrasive grinding has been an inability to effectively use grinding wheels with small diamond grit without having to constantly dress the wheel. This must be done to expose fresh sharp diamonds and remove the buildup of “swarf,” or ground glass, which binds between the exposed diamonds on the surface of the matrix.

The ELID process uses an electric current to enhance the surface oxidation of a conductive matrix holding the diamonds. This oxidation erodes the matrix, which prevents the build up of swarf and exposes fresh sharp diamonds, as illustrated in Figure 6.3-2. Therefore, as the wheel cuts the optic surface, it is also self-dressing and maintaining a good grinding surface. Since they are self-dressing, the ELID grinding-wheels' wear rate is significantly lower than that of conventional fixed-abrasive wheels, which require very aggressive dressing to remove swarf, expose fresh diamonds, and return their cutting surfaces to a nominally plano figure. Small-scale development results for this technology, as applied to phosphate-based laser glass, are illustrated in Table 6.3-2.^{4,5} Analysis of these data shows an improvement in material removal, a reduction in machine hours, and less wheel wear with the ELID technique. In addition, the resulting surfaces all had sufficiently good finishes to go directly to the polishing step. As this is being written, a large-scale demonstration of this technology is under way on equipment that will be used to manufacture large optics for NIF.

For a high-speed polish out, Tinsley developed a computer-controlled, high-speed grinder/polisher to rapidly lap (fine grind) lenses and flats and then polish them out (Figure 6.3-3). An advanced version of the overarm spindle machines commonly used in lens production, its advanced features include automated control of the part and tool speeds, and of the pressure between the tool and optic. The high-speed lapping tool is an array of bound abrasive pellets attached to a backing plate, which is matched to the shape of the surface being finished. The machine and process are now optimized to remove 500 μm of material from the surface of a typical 440-mm-square optic to a 9- μm finish, in about two hours.

Table 6.3-2. Summary of ELID vs conventional grinding performance data (time in hours).

Test condition	Coarse and fine generation time	Grinding time	Polish out time	Wheel wear rate ($\mu\text{m/hr}$)
ELID grinding	0.46 ± 0.04	0.66 ± 0.03	4.83 ± 0.69	11.0
Conventional grinding	0.17 ± 0.00	3.24 ± 0.47	5.67 ± 0.47	47.1





Figure 6.3-2. A comparison of traditional and ELID grinding.

By replacing this lapping tool with a similar tool covered with a synthetic lap material, the same machine can also be used for high-speed polishing. During a high-speed polish out of the same part, an additional 36- μm -thick, fused-silica layer can be removed in three to four hours.



Figure 6.3-3. The high-speed lapper/polisher developed by Tinsley for a high-speed polish out of diffuse scattering (“gray”) using pellet grinding.

Deterministic Figuring and Precision Polishing

All three vendors investigated methods of figure control using advanced, highly deterministic processes and modern optical metrology tools and control systems. Figure control for large optics is generally accomplished in one of two ways: precision planetary polishing or computer-controlled, small-tool polishing.

The primary impediments to the planetary polishing of large precision plano-optics are the high capitalization costs for large planetary polishers and the frequently encountered lack of determinism associated with this technology. Both Zygo and Kodak were very successful with combinations of machine improvements, new diagnostics, and metrology—not only in adding determinism to their planetary polishing processes, but also in obtaining projected throughput increases of three to six times and lowering costs for their planetary polishers (Figure 6.3-4).

The shape of the pitch-polishing lap has a significant impact on an optic’s final figure. At Zygo, a computer algorithm was devised that maintains a lap’s flatness to within $\lambda/20$ -wave, peak-to-valley, as measured on a 10-inch glass monitoring plate. Previously, without the intervention of a highly skilled optician, lap flatness diverged up to $\lambda/3$ wave, peak-to-valley, from a flat surface over the same 50-hour interval—as illustrated in Figure 6.3-5. This work has resulted in a computer-controlled, pitch-polishing lap that maintains a desirable figure, thus freeing the operator to concentrate on other tasks.




Figure 6.3-4. The 168-in. continuous planetary polisher used for the deterministic figuring of large NIF flats.




Figure 6.3-5. Without computer control or highly skilled operator intervention, the surface developed by a polishing lap diverges from flatness.

In the past, inherent wavefront ripple has limited the performance of optics fabricated by subaperture figuring, such as that done with a small-tool polisher. Ripple refers to periodic residual errors with a magnitude of 1 to 6 nm ($\lambda/600$ to $\lambda/100$) for spatial scalelengths of from 0.1 to 30 mm. In high-power laser applications, wavefront ripple causes dispersion that leads to near-field beam-intensity modulation. Wavefront ripple with spatial periods from 2 to 10 mm are nonlinearly amplified, which can result in laser damage and interference with beam propagation.^{6,7}

In an effort funded by LLNL, Eastman Kodak developed a process using a computer-controlled, small-tool polisher (Figure 6.3-6), by which large (flat, spherical, or aspheric) optics can be polished “ripple-free.” Kodak’s process developments resulted from advances in metrology, which were combined with Fourier analysis and extensive finishing-process modeling.




Figure 6.3-6. Kodak’s small-tool machine.

Specifically, Kodak developed an algorithm for assessing mid-spatial-frequency errors from phase measuring interferometry data, which employs Fourier analysis to identify the signature frequencies and orientations of the periodic errors. Then, by modeling the motion of various mechanical and electronic parameters across the aperture, they are able to match—and then reduce or eliminate—the source terms giving rise to the signature frequencies.

Figure 6.3-7 compares a typical subaperture wavefront with mid-spatial-frequency ripple, from a laser-amplifier slab—after conventional small-tool polishing (a)—to a similar wavefront produced from a sample polished using Kodak’s new small-tool polisher process (b). Using these data, the 1D plots of the power-spectral-density (PSD) analysis for the respective parts are compared in Figure 6.3-8. For reference, the line shown in each plot represents the PSD specification as a function of reciprocal mm (mm^{-1}) for the NIF laser. Figure 8 shows that Kodak significantly reduced the periodic ripple errors by an order of magnitude in PSD amplitude, while figuring a 0.9-m optic to better than $\lambda/6$ peak-to-valley flatness.




Figure 6.3-7. (a) The transmitted wavefront from a Beamlet amplifier slab with significant small-tool ripple. Note the grid-like striations with ~ 10 -mm spacing across the part. (b) Process improvements made by Kodak have significantly reduced variations in the transmitted wavefront from this slab.




Figure 6.3-8. A 1D power-spectral-density (PSD) analysis of the two slabs described in Figure 6.3-7. The earlier Beamlet slab (a) has a spatial scale ripple amplitude (at ~ 12 mm) that is 10 times higher than that of (b).

This “ripple-free,” computer-controlled, small-tool polishing process capitalizes on the advantages of other small-tool processes (in terms of high convergence and cost effectiveness), while also achieving a precision of figure and ripple control similar to that of the continuous-polishing process. For large-flat optics fabrication, significant cost savings in labor and capital investment result, compared to conventional flat-polishing techniques.

Tinsley has developed another small-tool technology, with similar capabilities, for figuring and polishing large flat, spherical, or aspheric optics. Their early efforts

to apply small-tool figuring to lenses suffered from ripple in the 6- to 12-mm spatial regime, as can be seen from the PSD of a Beamlet lens, shown in Figure 6.3-9a. LLNL funded work at Tinsley to reduce ripple power, with results similar to those demonstrated by Kodak on amplifier slabs. Ripple was reduced through proprietary process changes in machine precision, dwell times, by adding randomness, and by eliminating periodic-error functions. The improvement is apparent in Figure 6.3-9b, which shows the PSD for a NIF SF1/2 cavity spatial-filter lens fabricated at Tinsley in late 1998. This is just one of several Tinsley-fabricated NIF first-bundle lenses that meet all NIF specifications. High-resolution, reflected-wavefront measurements were made with a microinterferometer (Figure 6.3-10a) to evaluate the PSD-II (waviness-2) spatial periods (between 2.5 and 0.12 mm) and the NIF PSD-III roughness specification (over spatial periods > 0.12 mm), as shown in Figure 6.3-10b.




Figure 6.3-9. Comparison of power spectral density curves for (a) a Beamlet diagnostic lens with small-tool ripple vs (b) a NIF SF1/2 cavity spatial filter lens, where ripple was reduced to meet specifications for PSD-I over spatial periods of from 33 mm to 2.5 mm.




Figure 6.3-10. (a) NIF SF1/2 cavity spatial filter lens meeting the NIF PSD-II waviness-2 specification for spatial periods ranging from 33 mm to 0.12 mm and the PSD-III roughness specification for spatial periods > 0.12 mm (b).

Compensating for Material Inhomogeneities

Total system cost minimization has always been one of the primary goals of LLNL's optics-development efforts. One opportunity for cost minimization is the relation between material inhomogeneities and finishing processes. In the past, the fused-silica inhomogeneity specification for lenses and windows on large laser systems (such as for Nova and Beamlet) has been at two parts per million—the highest quality material that could reasonably be purchased from Corning or Heraeus in large quantities for the apertures of interest. The inhomogeneity is $\Delta n = e\lambda/t$, where e is the error in fractions of a wavelength, λ is the measurement wavelength, and t is the thickness of the part along the path traveled by the light. When measured at 633 nm, two parts per million translates into a transmitted wavefront error of $\sim\lambda/7$ for a 46-mm-thick NIF spatial-filter lens. Thus, at an inhomogeneity of 2 ppm, the material consumes about 45% of the transmitted wavefront specification of $\lambda/3$ for NIF spatial filter lenses. This leaves 55% of the transmitted-wavefront error budget for the finishing vendor.

LLNL has pursued aggressive cost targets for all steps in optics fabrication, including purchasing the fused-silica blanks. One lever used to achieve a lower cost for fused silica was recognizing that finishing processes can compensate for certain types of inhomogeneities in the blank. This tradeoff means buying less-expensive fused silica and taking advantage of the recently developed capabilities of our finishing vendors.

Analyses of typical material inhomogeneities, and our detailed knowledge about the finishing processes, have led to material specifications that give substantially higher yields (at lower costs to us) from the material vendor, without significantly increasing the cost of optic finishing. For instance, power (spherical errors) and astigmatism (cylindrical errors) are typical residual errors that are corrected during the course of continuous planetary polishing by the optician during final polishing. The techniques used by opticians to do this include weighting and controlling the lap's shape through the position of the conditioner. The process used to correct a residual error is the same, whether the error is due to irregularities in the surface figure or to inhomogeneities in the material. As long as the errors vary slowly over the surface and the optic is tested in transmission while it is being corrected, material inhomogeneities can be compensated for by a surface-figure change. For instance, the NIF window-blank specification permits up to $\lambda/4$ of power to be subtracted from the homogeneity error before the final analysis. Indeed, a $\lambda/4$ -power error is typically encountered and removed during final finishing.

As just described, both Tinsley and Kodak developed small-tool figuring processes that can remove transmitted wavefront errors with scale lengths larger than the size of the polishing tool. All the lenses in the NIF are aspheric, with small spherical surface-figure departures of between 1.4 and 9 μm . Because of these small aspheric departures, the NIF lenses are typically processed through all of the conventional steps for producing a spherical lens; then the required aspheric departure from the spherical surface is polished in. Since polishing can only remove material, the correct surface form is approached in a step-wise manner by an iterative process. Each

polishing step is followed by a precision-metrology step, which identifies where (and how much) additional material must be removed. If too much material is removed from the wrong location, the whole surface must be brought down to bring the final surface to the correct form. To be conservative, the first polish of the aspheric zones of the lens is necessarily shallow and the residual error can be on the order of several waves.

To take advantage of the fact that small-tool figuring is typically used to produce the NIF's lenses, the material specification for lens blanks was relaxed to allow subtracting up to $\lambda/4$ of spherically symmetric error (power) and up to $\lambda/7$ of nonsymmetric long-wavelength error, before calculating the material's inhomogeneity. Long-wavelength errors were defined as errors with a spatial period >100 mm—significantly greater than the typical tool size.

Figure 6.3-11 shows a fused-silica blank, especially selected to verify that long-wavelength errors of this type and magnitude can be removed by small-tool figuring without significantly modifying the standard lens fabrication process already developed at Tinsley for NIF's spatial-filter lenses. In addition to the long-wavelength error that we intended to evaluate in this experiment, this particular blank also had a small diameter (30-mm) localized index inhomogeneity of $\sim\lambda/3$ amplitude (readily visible in Figure 6.3-11). This was associated with a collection of small bubbles, which was sliced off when the blank was removed from the boule. This type of localized inhomogeneity does not pass the NIF specification for P-V amplitude or for gradient rms ($<\lambda/100$ waves/cm) due to the high local gradients associated with the defect. (The blank was used despite the defect because it was difficult to obtain glass with long-wavelength error near the top end of the newly revised NIF lens-blank inhomogeneity specification.)




Figure 6.3-11. Transmitted wavefront from a fused-silica blank (oil-on-flat inhomogeneity) with a high long-spatial-scale inhomogeneity error (0.4λ). This blank also had a small-diameter, (~ 30 -mm) discrete inhomogeneity (0.35λ), associated with a cluster of bubbles that had been sliced off.

We expected that Tinsley's small-tool figuring process would remove long-wavelength errors in the blank, while most of the error associated with the localized inhomogeneity would remain. If the localized defect is ignored, the blank's index inhomogeneity was 5.4 ppm before the long-wavelength errors (those with spatial periods >100 mm) were subtracted—and 1.5 ppm after their subtraction. With the localized defect included, the blank had an inhomogeneity of 8.4 ppm, which far exceeds the 2-ppm historical specification for optics of this type. For a NIF spatial filter lens, this level of inhomogeneity translates to a 0.6-wave error—twice the NIF specification for P-V error on the finished lens.

Figure 6.3-12 shows the lens after the first cut at polishing the aspheric surface form into the lens. The total transmitted wavefront error at this early stage in figuring is 2.8λ . The long-wavelength material errors from the fused-silica blank are insignificant, compared to the residual errors resulting from the normal processing of the part through small-tool aspherization. This suggests that it is possible to relax the amplitude specification for allowable long-wavelength errors in lens blanks even further than $\lambda/7$, without affecting the quality of the finished NIF lenses. After the final small-tool processing, the total transmitted wavefront error for this SF1/2 cavity spatial-filter lens was reduced to $<\lambda/5$ vs the specification of $\lambda/3$ —even including the error associated with the localized inhomogeneity.




Figure 6.3-12. The transmitted wavefront from a lens that has been partially fabricated from the fused-silica blank shown in Figure 6.3-11. The discrete inhomogeneity is readily apparent, but the long spatial-scale error in the raw material is insignificant, relative to the total amount of error induced at this early stage of the work.

The localized defect is apparent in Figure 6.3-13, but it was reduced in amplitude by a factor of seven during normal small-tool figuring to $<\lambda/20$. The finished lens met all NIF PSD and wavefront specifications, except in the subaperture containing the residual error from the original discrete inhomogeneity. This experiment

demonstrated the ease with which small-tool figuring of NIF lenses can accommodate a significant amount of long-wavelength error. It also reinforced the importance of avoiding localized inhomogeneities in the raw material that are too small in scale to be removed by the small tool.

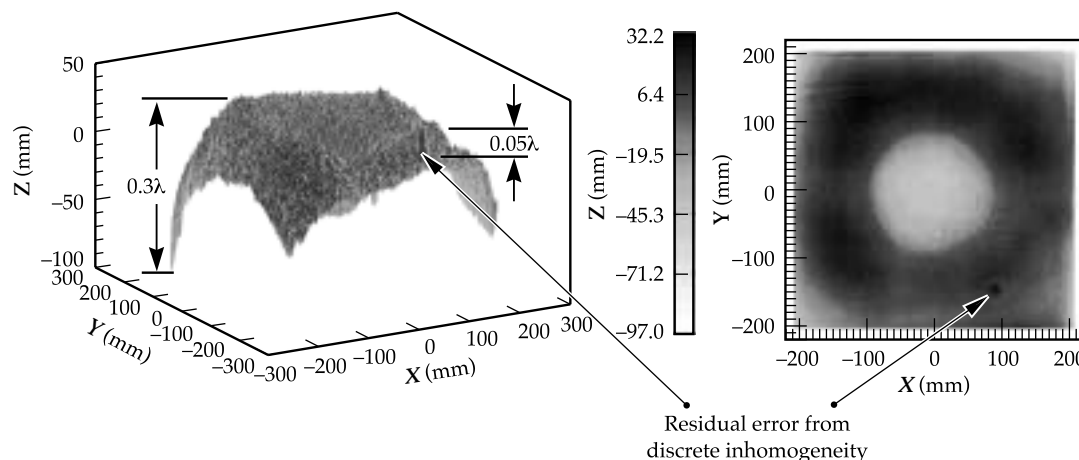


Figure 6.3-13. The transmitted wavefront from the finished NIF cavity spatial-filter lens made from the blank shown in Figure 11. The discrete inhomogeneity is still visible, but its amplitude has been reduced seven fold.

Lenses are the only NIF optics with small-tool correction of long-wavelength material inhomogeneities built into the specifications for the blanks. However, laser glass does have a specification allowing long-wavelength material inhomogeneities including power (0.425λ) and astigmatism (0.220λ), which are spherical and cylindrical terms, respectively. These simple long-wavelength features are correctable on plano optics with either a small tool or a pitch lap. With a pitch lap, the machine has to be operated in a condition that will yield intentionally nonflat optical surfaces, rather than the flat surfaces normally desired.

To cancel the material-induced wavefront distortion, a surface of the inverse sign and equal magnitude must be polished into the optic. This technique can also be used for fused-silica windows or BK7 polarizers, if required. In some cases, the spatial period of the material inhomogeneity might be too short to correct on a pitch lap, requiring the use of a small-tool figuring step in addition to the traditional flat-finishing manufacturing process.

Finishing for a High 3ω Laser-Damage Threshold

Finishing-development efforts, both at LLNL and at several vendors, have demonstrated that there are several paths toward obtaining fused-silica optics with finished surfaces meeting the NIF's 3ω fluence requirements: an average fluence of up to 8 J/cm^2 with peak fluences about 1.5 times higher.

For high-quality, fused-silica substrates, it is generally accepted that the damage performance of the polished surfaces is largely controlled by micron-scale subsurface damage, or by nanoscale absorbing particles. The subsurface damage results from flaws in the material generated by the various preliminary grinding and polishing steps.⁸ The absorbing defects are due to contamination in the nominally 100-nm-thick polishing redeposition layer (Figure 6.3-14).




Figure 6.3-14. A conceptual illustration of subsurface damage.

Several rules have been developed which, when followed, help to minimize subsurface damage and redeposition layer contamination. Following these rules has led to a general improvement of the laser-damage performance of the polished surfaces. They can be summarized as falling into two main groups:

1. **Subsurface damage.** High-performance parts are created through a series of finishing steps, each one of which uses successively smaller grinding particles. The depth of material removed in each step should be about three times the maximum size of the polishing particles used in the previous step. Knowing the size distribution of the particles used at each stage is necessary, as the maximum particle size is often two to three times larger in diameter than the average particle size.⁹ Figure 6.3-15 shows the effect of subsurface damage on the laser-damage threshold for three parts with various degrees of residual subsurface damage.¹⁰
2. **Polishing contamination.** The chemomechanical nature of the final polishing operations on fused silica results in a smooth layer of silica being deposited, which is contaminated with the range of elements present in the polishing slurry. (These include cerium [Ce] from the ceria [CeO₂] polishing compound, plus iron, chromium, copper, sodium, potassium, and aluminum, which are present in the slurry either as anti-agglomerants or are incorporated into it as contaminants during the polishing operation.) The effects of polishing contamination can be minimized by selecting nonabsorbent polishing compounds, or by removing the contaminated re-deposition layer.

Depth profiling by secondary-ion mass spectroscopy (SIMS) shows that contaminants are present in the ~100-nm-thick redeposition layer, at levels from 1 to 1000 ppm. Figure 6.3-16 illustrates a concentration profile of Ce in the polish layer of a ceria-polished, fused-silica optic.¹¹ Ceria, a widely used polishing compound, is a strong absorber at 351 nm ($\alpha \approx 70,000 \text{ cm}^{-1}$). Simple thermo-mechanical models indicate that a 3-ns, 15 J/cm^2 laser pulse could heat a 10-nm-diameter ceria sphere to ~4000 K, therefore leading to a thermal explosion.¹² Fe, Cr, and Cu oxides would also be heated to similar levels.




Figure 6.3-15. The level of grind and polish was varied on these three ceria-polished unetched parts—to demonstrate that the first step toward increasing the damage threshold is subsurface damage control.




Figure 6.3-16. SIMS depth profile of Ce concentration inside the redeposition layer of a polished fused-silica surface.¹¹

The strongest indication of contaminant-related damage at high fluences is the “gray haze” morphology present on damaged ceria-polished surfaces. This haze consists of 1- μm -diameter damage pits, with concentrations of from 10^3 to $10^5/\text{cm}^2$ (Figure 6.3-17).¹³




Figure 6.3-17. A SEM micrograph of gray-haze damage created by the 351-nm, 7-ns illumination of a polished fused-silica surface. ¹³

The gray-haze damage can be reduced or eliminated by at least three different approaches:

1. Chemically etching the polished surface to remove contaminants. The gray-haze damage is attributed to ceria contamination in the 100-nm-thick redeposition layer. If this layer were to be removed with a hydrofluoric acid (HF) etch, it is expected that the damage morphology would also be removed. Studies of damage performance as a function of etch depth found that for 300 to 600-nm¹⁷ deep etches, the gray haze was significantly reduced. Figure 6.3-18 compares the fluence vs damage threshold behavior of two parts similarly polished with ceria.¹⁰ The HF-etched part exhibited no gray haze and performed significantly better in the high-fluence portion of the testing. Figure 6.3-19 shows a ceria-polished Beamlet debris shield, which was HF etched over one-third of its area to a depth of 600 nm.¹⁷ For the as-received unetched part, high concentrations of gray-haze damage sites were observed for fluences above $12 \text{ J}/\text{cm}^2$. After the 600-nm-deep etch of an adjacent area of the part, the damage concentration was significantly reduced. Typically, etching the surface layer also exposes finishing flaws hidden by the redeposition layer; these are readily apparent in Figure 6.3-19.
2. Removing the contaminated surface layer with an ion etch. An alternative to a wet etch could be an ion etch of the surface layer. The Laboratory for Laser Energetics (LLE) has reported that the damage thresholds of phase-plate substrates are generally improved after an ion-beam etching.¹⁴ Researchers at Osaka University have reported that after an argon-ion-beam etching of subsurface damage, the 355-nm damage threshold of fused-silica surfaces was doubled.¹⁵ Ion-etch studies at LLNL show that there is a higher chance of

avoiding gray-haze damage with an ion etch than without one. However, the laser-damage data obtained from these ion-etched surfaces did not conclusively show that ion etching improved the local damage performance for the parts studied.




Figure 6.3-18. Etching with HF generally improves the high-fluence performance of ceria-polished optics by eliminating the gray-haze damage morphology.¹⁰

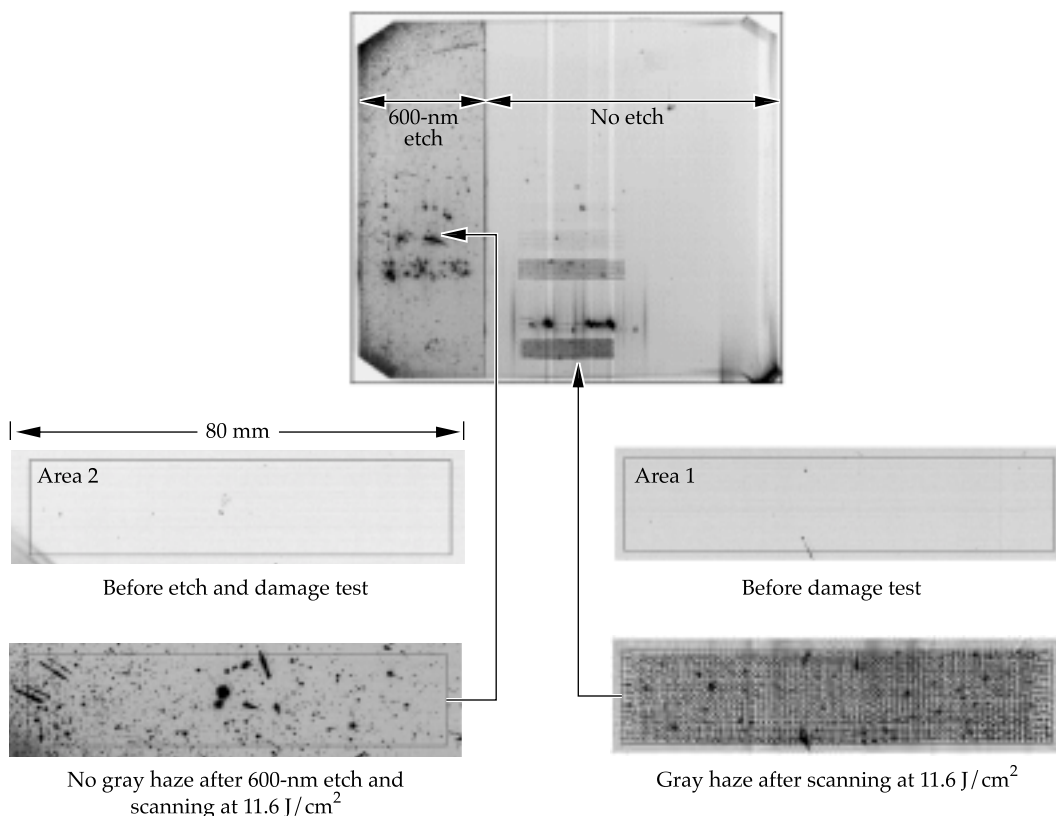


Figure 6.3-19. Gray haze can be eliminated from ceria-polished optics (or reduced) by a 200- to 600-nm-deep HF etch to remove the contaminated polish layer. In general, the etching slightly increases scattering because it uncovers hidden scratches and other subsurface damage.¹⁶

3. Final polishing with zirconia instead of ceria. Zirconia has a negligible absorption at 351 nm and is therefore less likely than ceria to be a damage precursor. Surfaces polished with zirconia consistently do not show gray haze damage, while ceria-polished surfaces almost always show some level of gray haze at fluences near 10 J/cm^2 , with 3-ns pulses. Figure 6.3-20 compares two fused-silica samples that were similarly polished using zirconia on one and ceria on the other.¹⁰ The zirconia polished surfaces show higher R:1 damage distributions than the ceria polished part. However, great care must be taken in choosing slurry parameters for zirconia polishing, as the material is easily agglomerated, resulting in surface scratches. The scratches show various degrees of propensity towards damage. Damage thresholds have been observed to decrease with zirconia polishing—if subsurface damage and scratch densities are not allowed to increase, due to poor slurry-parameter control.




Figure 6.3-20. As long as subsurface damage is properly managed, the damage threshold can be improved by polishing with zirconia instead of ceria.

Recent research at LLNL, at our finishing vendors, and elsewhere has led to new developments demonstrating that fused-silica optics can be fabricated to meet the NIF's 3ω fluence requirements. Meeting these targets requires combining careful attention to the classic techniques of grinding and polishing (for controlling subsurface damage) with an understanding of the role played by contaminants in the polish redeposition layer.

Despite the excellent progress made in increasing damage thresholds and reducing damage density, it is clear that some damage will still occur at the NIF's operating fluences. Understanding how these more damage-resistant surfaces will behave in the NIF requires defining the statistical relations of damage density and damage growth with laser fluence and shot history.¹⁷

Notes and References

1. L. J. Atherton, "Optics Manufacturing Readiness," *Manufacturing Readiness Plan for the National Ignition Facility*, Lawrence Livermore National Laboratory, Livermore, CA, NIF-LLNL-94-204, pp. 9–55, April 1994.
2. D. M. Aikens, L. Rich, D. Bajuk, and A. Slomba, "Developing Enabling Optics Finishing Technologies for the National Ignition Facility," Lawrence Livermore National Laboratory, Livermore, CA, UCRL-JC-129317, 1998 Summer Topical Meeting, Kailua-Kona, HI, June 8–12, 1998.
3. H. Ohmori, *Int. Jour. Japan Soc. Prec. Engr.* Vol. **26**, 273–278, (1992).
4. K. Grobsky and D. Johnson (private communication), 1995.
5. F. Tinker and J. Hopkins (private communication), 1997.
6. C. R. Wolfe, J. K. Lawson, M. C. Kellam, R. T. Maney, and A. Demiris, "Measurements of Wavefront Structure from Large-aperture Optical Components by Phase Shifting Interferometry," in *Optical Manufacturing and Testing*, Victor J. Doherty, H. Phillip Stahl, eds., in *Proc. SPIE* **2536**, 13–37 (1995).
7. J. K. Lawson, C. R. Wolfe, K. R. Manes, J. B. Trenholme, D. M. Aikens, and R. E. English, Jr., "Specification of Optical Components Using the Power Spectral Density Function," in *Optical Manufacturing and Testing*, Victor J. Doherty, H. Phillip Stahl, eds., in *Proc. SPIE* **2536**, 38–50 (1995).
8. P. P. Hed, D. F. Edwards, and J. B. Davis, *Subsurface Damage in Optical Materials: Origin, Measurement and Removal*, Lawrence Livermore National Laboratory, Livermore, CA, UCRL-99548 Rev 1 (1998). Submitted to *Appl. Opt.*
9. M. Nichols et al., LLNL Patent Docket No. IL-10215, (1998), *Fabrication of Optical Component*.
10. D. W. Camp et al., *Subsurface Damage and Polishing Compound Affect the 355-nm Laser Damage Threshold of Fused Silica Surfaces*, Lawrence Livermore National Laboratory, Livermore, CA, UCRL-JC-129301. Prepared for *Laser-Induced Damage in Optical Materials: 1997*, in *Proc. SPIE* **3244**, 356–364 (1998).
11. M. Kozlowski et al., "Depth profiling of Polishing-induced Contamination on Fused Silica Surfaces", *Laser-Induced Damage in Optical Materials: 1997*, in *Proc. SPIE* **3244**, 365–375 (1998).
12. M. D. Feit et al., "Modeling of Laser-Induced Surface Cracks in Silica at 355 nm," *Laser-Induced Damage in Optical Materials: 1997*, in *Proc. SPIE* **3244**, 350–355 (1998).
13. J. Yoshiyama et al., "A Study of the Effects of Polishing, Etching, Cleaving, and Water Leaching on the UV Laser Damage of Fused Silica," *Laser-Induced Damage in Optical Materials: 1997*, in *Proc. SPIE*, **3244**, 331–340 (1998).
14. D. J. Smith, J. A. Warner, N. LeBarron, and S. LaDelia, "Production of Distribution Phase Plates Using an Energetic Ion Process," *Annual Symposium on Optical Materials for High Power Lasers*, Sept. 28–Oct. 1, 1998.
15. T. Kamimura et al., "Improvement of Laser-Induced Surface Damage in UV Optics by Ion Beam Etching," *Annual Symposium on Optical Materials for High Power Lasers*, Sept. 28–Oct. 1, 1998.
16. L. Sheehan and R. Moaser, unpublished results, Lawrence Livermore National Laboratory, Livermore, CA (1998).

17. M. D. Feit et al., "Extrapolation of Damage Test Data to Predict Performance of Large-Area NIF Optics at 355 nm," *30th Annual Symposium on Optical Materials for High Power Lasers*, Boulder, CO, Sept. 28–Oct. 1, 1998.

6.4 NIF Optic Processing

Jim Fair, Chuck Thorsness, Pam Whitman

6.4.1 Introduction

The LLNL Optics Processing and Development Laboratory (OPDL) will provide cleaning, coating, and subassembly services for beamline optics during the lifetime of the NIF. Processes and equipment used in the NIF OPDL are the result of many years of experience handling Nova optics, adaptation of processes used at the Laboratory for Laser Energetics (LLE) OMEGA laser facility, and new developments tailored to meet specific NIF requirements. During 1999, the conversion of 10,000 ft² of Nova building space into a clean manufacturing floor for the OPDL was completed. Most of the processing equipment for OPDL operations has been installed and process validation using pilot optics and prototype fixtures is under way.

The NIF requirements for stable, high-performance antireflection coatings have necessitated improvements in existing sol-base antireflection coatings. During 1999, significant progress was made in understanding the relationship between sol film porosity and the tendency of the films to adsorb organic contaminants from the atmosphere. A post-deposition ammonia vapor treatment of the sol films was shown to significantly reduce the coating affinity for organic contaminants, primarily through removal of interparticle microporosity. The elimination of microporosity produces a film that shows 100% reversibility of organic adsorption. This property presents the opportunity to use a sacrificial adsorbent to protect the optic coatings from contamination during service.

6.4.2 Optics Processing Facility

The optics processing laboratory is divided into multiple work areas that permit appropriate transitions from the relatively dirty external environment to the internal Class-100 processing areas (Figure 6.4-1).

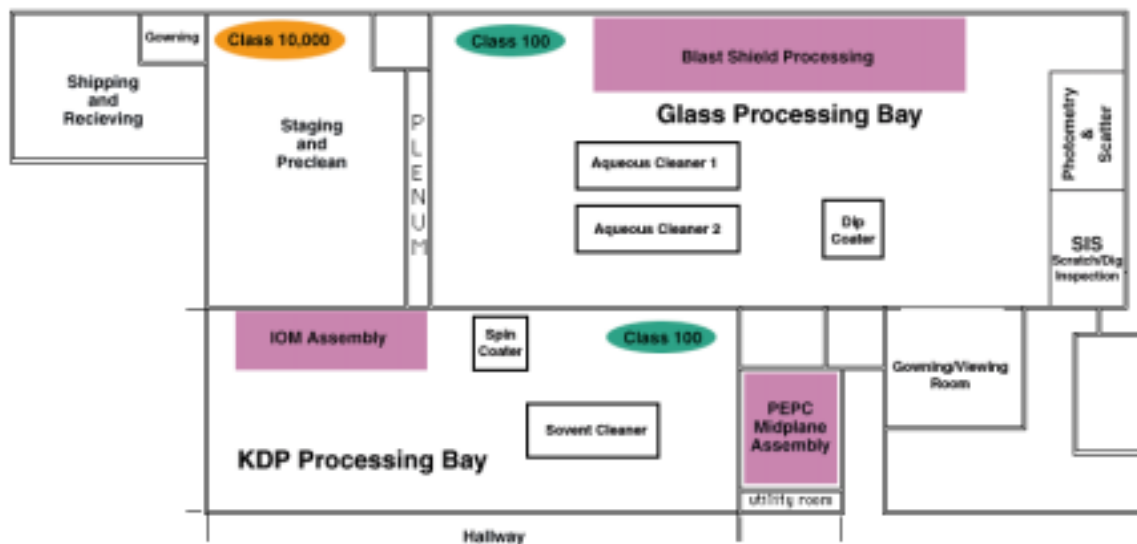


Figure 6.4-1. OPDL laboratory layout.

The physical isolation between glass and KDP handling areas is necessitated by the high sensitivity of KDP to trace levels of ammonia potentially present as a normal consequence of dip coating and ammonia curing operations on the glass processing floor. Airborn particulate levels are monitored continuously throughout the facility by an automated multi-port sampling system (Figure 6.4-2). As Figure 6.4-3 shows, the entire manufacturing floor typically operates well below Class 100.




Figure 6.4-2. History of air born particulate levels in OPDL.

Table 6.4-1 summarizes OPDL services and shows the interfaces to other key NIF facilities and organizations. Although current activities are focused on establishing core capabilities for initial construction of the NIF, the OPDL has been designed to meet the long-term optic refurbishment requirements, which are projected to be as much as 20 times that of NOVA.

Table 6.4-1. Summary of OPDL services to the NIF.

Optic type	Cleaning	Sol AR coating	Coating cure	Sub-assembly	Customer
Amplifiers, Mirrors, Polarizers	Aqueous/ultrasonic	n.a.	n.a.	no	OAB
Deformable mirror	Aqueous/manual	n.a.	n.a.	no	OAB
Spatial filter lenses, PEPC switch windows, Diagnostic beam splitter (1 ω), 1 ω phase plate	Aqueous/ultrasonic	1 ω quarter-wave sol-dip	NH ₃ /H ₂ O vapor treatment	Yes. Spatial filter/beam splitter subassembly and PEPC midplane	OAB
3 ω focus lens, Debris shields, 3 ω sampling grating	Aqueous/ultrasonic	3 ω quarter-wave sol-dip	NH ₃ /H ₂ O vapor treatment	Yes. Debris shield and grating cassettes, IOM	Target chamber
PEPC switch Crystals	Toluene/ultrasonic	1 ω quarter-wave-spin	none	Yes. PEPC midplane	OAB
KDP Frequency Doubler	Toluene/ultrasonic	1 ω quarter-wave, 1-2 ω quarter-wave, silicone/sol 2-layer spin	none	Yes. Final optics cassette	Target chamber
DKDP frequency tripler	Toluene/ultrasonic	1-2 ω quarter-wave, sol/silicone-spin 3 ω sol-spin	none	Yes. Final optics cassette	Target chamber

Cleaning and Coating Equipment

Fixtures and Handling. The dimensional similarity of large subsets of NIF optics presents the opportunity to use universal fixtures and handling interfaces. Finished KDP crystals are held by a single fixture (Figure 4-6 in the FOA section) during shipment and all subsequent processing steps. Although contact pads of appropriate material are substituted at various intermediate steps, the crystalline optic remains with

a single frame from finishing to final insertion of the crystal into the NIF optic cassette or Pockels cell midplane.

Amplifier slabs, mirrors, and polarizers require precision cleaning before shipment to the OAB for insertion into line replaceable units (LRUs). A universal wash frame (Figure 6.4-3) with movable contact blocks has been designed to accommodate multiple optic types with only minimal adjustment. The fixture has no wetted overhead parts and, therefore, does not interfere with uniform water sheeting and even drying of the cleaned optic surface.



Figure 6.4-3. Large optic wash frame held by Alupalift™-based handling tool.

Optics that require a sol antireflection coating are held by a dual function cleaning/coating frame (Figure 6.4-4) during processing. As with the wash frame, the

handling interface points of the coating fixture are designed for compatibility with the standard lifting and transport tools.



Figure 6.4-4. Large optic cleaning and coating frame being loaded into dip-coating system.

A battery-powered lift truck with a universal fork interface is used to transport optics between processing points within the facility (Figures 6.4-3 and 6.4-4). A similar custom optic grasping tool is used to transfer optics to and from the cleaning fixtures, transport containers, and metrology equipment (Figure 6.4-5).




Figure 6.4-5. Alupalift-based universal handling tool loading laser glass slab into laser surface scanning tool.

Cleaning Systems

Glass Cleaning. Hand cleaning with water, soap, mild abrasives, and alcohol has been the traditional method for the cleaning of large glass laser optics prior to coating and/or assembly. The automated cleaning systems for the NIF (Figure 6.4-6) have been designed to accomplish the same cleanliness level with ultrasonic scrubbing in aqueous solutions. The baseline process for laser glass cleaning has been established, and during FY00, the fused silica processes will be finalized.



Figure 6.4-6. JST aqueous/ultrasonic cleaning systems.

KDP Cleaning. The NIF cleaning process for KDP-KD*P crystals must consistently remove the remnants of oils used during final diamond tool finishing of the optics, the KDP debris that results from the diamond flycutting process, and any particulate contamination that results from handling the optics prior to the application of the antireflection coating. Two identical cleaning systems have been built by Forward Technology Industries, Inc. for ultrasonic solvent cleaning of KDP crystals. One system (shown in Figure 6.4-7) has been installed, commissioned, and is now operating at LLNL for final precision cleaning of all NIF crystals. The second system will be installed at the crystal finishing vendor.



Figure 6.4-7. Forward Technologies KDP cleaning system.

Cleaning process development focused in two areas this past year: (1) identification of appropriate solvents to remove the organic diamond finishing oils and (2) identification of process parameters and solvents to remove KDP debris and ambient particulate contamination. A sequence of more polar to less polar solvents (between 2-butanol, methyl isobutyl ketone (MIBK), and toluene) was tested for both organic contamination and particulate removal from freshly diamond-turned KDP surfaces. So far, toluene has been the most effective solvent for removal of likely diamond turning contaminants.

Coating Systems

Glass Coating. Dip coating has been the mainstay for application of antireflection films on large laser optics for many years at LLNL and LLE. The new NIF dip coat system (Figure 6.4-8) was patterned after a similar system at LLE in Rochester, NY, which uses a ridged mount and lead-screw drive instead of a free suspension cable.




Figure 6.4-8. Chemat dip coat system.

For fire protection reasons it was necessary to surround the coating system by an airtight enclosure to provide vapor isolation from the laboratory environment. Sol coatings for the large fused-silica optics are deposited from an ethanol suspension of 3% silica. At this suspension level, typical coating rates are in the 5–15 cm/min. range. The total coating time for an optic is typically 30 to 45 minutes, with loading, unloading, settling times, etc., included. The sol suspensions are manufactured from high-purity tetraethylorthosilicate here at LLNL.

KDP Coating. The performance requirements for NIF conversion crystals necessitate depositing a different coating thickness on each surface of the crystal to match the use wavelengths. To support these requirements, a commercial spin coater (Figure 6.4-9) was purchased for coating all of the KDP and DKDP crystals. The spin coater can automatically dispense up to three different coating compositions, allowing the flexibility to separately optimize coatings for surfaces that must transmit 1ω , 1ω and 2ω , or 3ω laser light.



Figure 6.4-9. Suss spin-coating system.


Sol Coating Improvements

Organic contamination of the porous antireflection coatings on laser optics can be a significant contributor to loss of light transmission during service. Much of the affinity for organic contaminants originates in the large amount of microporosity within the 10- to 20-nm-diameter constituent silica particles of the sol coatings. Treatment of the sol coatings with aqueous ammonia vapor decreases the specific surface area (Table 6.4-2), largely by the removal of microporosity.

Table 6.4-2. Effect of vapor treatment on specific area and pore size in bulk and film sol samples.

Sample	Specific Area (m ² /g)	Pore Peaks (Å)
As Deposited Film	174	70
NH ₃ Treated Film	76	135 & 18
Evaporated Bulk Sol	564	55
NH ₃ Treated Bulk	238	100
NH ₃ +HMDS Treated Bulk	215	100

A further treatment of the coating with HMDS (hexamethyldisilazane) can further reduce the surface affinity for polar organics by coating a significant portion of the silica surface with nonpolar methyl groups. The ammonia treatment alone, also significantly increases the abrasion resistance of the coatings.¹ Figure 6.4-10 shows the effect of ammonia and HMDS treatment on the adsorption characteristics of DBP (dibutyl phthalate M.W. 278).

**Figure 6.4-10. Desorption isotherms at 20°C for DBP on porous silica coatings.**

The presence of high microporosity in the untreated coating is evident as a steep rise in adsorption near the origin accompanied by low pressure hysteresis, which accounts for an approximately 20% irreversible adsorption level when the untreated films are exposed to only trace amounts ($p/p_v < 0.01$) of DBP. In contrast, treated coatings show relatively little adsorption at low organic partial pressures and, therefore, admit some realistic operating margin in the presence of organics without unacceptable reflection loss (Figure 6.4-11).




Figure 6.4-11. Reflectance loss versus relative pressure of DPB for ammonia treated, 1 ω sol antireflection coating.

Capillary condensation in the larger mesoporous regions between silica particles dominates at higher partial pressures, limiting the allowable partial pressure of organics to approximately 0.2 of its pure component vapor pressure for treated films. This 0.2 P_v limit has been adopted as the maximum allowable organic level inside sol-coated optic enclosures for the NIF.²

Protection of Sol Coatings during Service

In addition to chemical treatment of the antireflection coatings, a promising approach to protection of porous antireflection coatings during service is to supply an overwhelming amount of sacrificial surface area inside the optic enclosure. A modest amount of activated charcoal can reduce the partial pressure of an organic contaminant to very low levels (Figure 6.4-12). Five hundred grams of activated carbon has the capacity to adsorb a minimum of 100 gm of organic contaminants and would, therefore, theoretically have enough capacity to remove all the organic contaminants from 5000 1 ω antireflection sol-coated NIF-size optics whose coatings were completely filled with organic contamination.




Figure 6.4-12. Isotherms for organic loading on activated carbon at 25 °C (from Calgon Carbon Corporation).

The removal of organic contaminants from a sol coating by activated charcoal has been demonstrated on a small scale using a surface-acoustic-wave device that measures directly the weight of adsorbed organic compounds within the sol coating. In Figure 6.4-13, the removal of DPB from a hardened sol coating by activated charcoal is shown. Notice that the fill fraction is driven from full saturation to essentially zero in less than two hours.




Figure 6.4-13. Removal of DBP from sol coating using activated charcoal at 5 mtorr.

Prediction of the degree of contamination of a sol coating and the expected practical gettering rate by activated charcoal in an optic enclosure is in general a complex question involving mass transfer, equilibrium, and geometric constraints. However, in the case when the sol can be considered to be in equilibrium with the adjacent vapor phase, it is relatively straightforward to relate contaminant levels in the vapor phase to the amount of contaminant loading in the sol by the use of an isotherm.

In many cases, the time to reach an equilibrium condition is long, or the equilibrium condition entails the net transport of material within a system. In these cases, the equilibrium considerations cannot be directly applied to determine the level of coating contamination at times of interest. To address this problem, some calculation of transport rates of contaminants within the system must be utilized.

We have developed a simple model to examine the rate processes occurring in sol contamination. The model is based on mass transfer considerations in the void volume of a system. Vapor contaminants in the void volume interact with surfaces characterized by adsorption isotherms such as those shown in Figure 6.4-10.

To model contaminant loading of the integrated optic module (IOM), mass transport is handled in an approximate fashion by breaking the mechanical system into regions. The vapor phase in each region is assumed to be in equilibrium with the wall in that region. Transport rates in the vapor phase are estimated by defining a conductance from region to region. The equilibrium partial pressure of contaminants in each region is used to define the mass transfer driving force and the transport rates. In generating the mass transfer conductance, use is made of Bosanquit's equation³ to obtain effective

diffusivities that span the operating regimes in which normal molecular diffusivity dominates to regimes in which diffusion is controlled by Knudsen (low pressure) flow.

The operating conditions of the IOM are not currently well established and may range anywhere from 10^{-4} Torr (0.013 Pa) to 10 Torr (1333 Pa). Contaminants of primary concern in the IOM tend to be things like machine oils, vacuum pump oils, and high-boiling-point plasticizers. To simulate these components, we have chosen to use DBP and dioctyl phthalate (DOP) as model compounds. Both are plasticizers and have vapor pressures at room temperature that span the range of interest, 1.5×10^{-5} T (2×10^{-3} Pa) and 1.5×10^{-7} T (2×10^{-5} Pa) respectively. These species, especially DOP, are often found as residue on many potential NIF components. Also, DBP is a reasonable model for many vacuum pump oils that have a similar vapor pressure.

Calculation results for coating contamination rates in an initially dirty IOM are presented below. These calculations are primarily aimed at exploring the potential usefulness of an activated carbon getter in protecting optical coatings in the IOM. In this analysis DBP and DOP are assumed to have identical normalized adsorption isotherms similar to that shown in Figure 6.4-10 (ammonia hardened case). The modeled IOM has been broken into 35 computational regions. Each region is associated with a metal, Spectralon, or optical-coating surface. Twelve coated optical surfaces are included: one on the vacuum window (VW) and the debris shield (DS) and two on the doubler crystal (CD), tripler crystal (CT), focus lens (FL), diffractive optics plate 1 (DO1), and diffractive optics plate 2 (DO2). Coating thickness characteristic of a 1ω antireflection are assumed to be present on the VW and on the front side of the CD. Two- ω coatings are assumed to be present on the backside of the CD and on the front side of the CT. All other sol coatings are assumed to have 3ω quarter-wave thickness. The activated carbon is assumed to be contaminated in a canister connected to the unobstructed front region of the IOM by a 3.8-cm (1.5-in.)-diameter tube 50 cm long.

Figure 6.4-14 shows the results of a calculation done at an operating pressure of 10^{-4} Torr. Fill fractions of each coated surface are presented. In this calculation, it is assumed that the coatings are originally clean and that the source of the contamination is from nonoptical surfaces in the IOM. An A-level of contamination (10^{-5} kg/m²) composed of 50% DBP and 50% DOP is assumed.




Figure 6.4-14. Computed fill fractions for each optic surface at an operating pressure of 10^{-4} Torr from level-A contamination of surfaces.

The rise and fall of contamination levels is a consequence of transport within the unit. Initially the contamination level rises as wall contaminants adsorb onto the coatings. Eventually, however, the levels fall as the activated carbon canister lowers the total contamination level within the IOM. Notice that the coatings on the internal optics of the final optics cell are considerably slower to respond than the other surfaces (note, A is the front side, and B is the backside of an optic). This is caused by the limited access to these surfaces from the bulk of the IOM walls. For the more exposed optics the 3ω (thinner) coatings tend to fill more rapidly because they require less total transport of material.

Figure 6.4-15 shows the effect of predicted contamination levels on overall transmission loss of the IOM. The shape of the computed transmission loss curve is a function of not only the level of contamination present, but also the shape of isotherms associated with the contaminated surfaces.




Figure 6.4-15. Computed transmission loss from coating reflectance changes in the IOM at an operating pressure of 10^{-4} Torr from level-A contamination of surfaces.

For a given level of contamination, the ability of the activated carbon to protect the optics is a function of operating pressure. Three calculations were done to show the influence of pressure. The computed transmission loss at pressures that cover the anticipated IOM operating range are shown in Figure 6.4-16. Notice that at higher pressures, the peak in the transmission loss moves to a later time because of reduced transport rates. At the highest pressure, the transmission loss would continue to climb during the entire period of projected operation. For operation at higher pressures where normal molecular diffusion plays a dominant role, the use of a clean purge flow would probably be preferred over the use of activated carbon in a quiescent system. The activated carbon (as a flow-through filter) could still play a role in this scenario as a means of locally insuring that the sweep gas entering each IOM contains no significant levels of organic contamination.




Figure 6.4-16. Computed transmission loss from coating reflectance changes in the IOM at three operating pressures from level-A contamination of surfaces.

With the data and modeling tools at hand, we are in a position to estimate the response of the IOM and other optic enclosures to various types and levels of organic contamination and to estimate the effects and interactions of possible mitigation strategies and operating conditions. While the approach taken here is appropriate for the indicated operating conditions, in situations where significant bulk flow mechanisms are active, more sophisticated modeling techniques would be required. Also, more basis data, in the form of isotherms for different surface and contaminant types, is needed to further substantiate this type of analysis.

References

1. P. F. Belleville and H. G. Flock, "Ammonia Hardening of Porous Silica Anti-Reflective Coatings," *Sol-Gel Optics III*, in Proc. SPIE (J. D. MacKenzie, ed.), Vol. **2288**, pp. 25–32 (1994).
2. J. Fair and C. Thorsness, "Maximum acceptable organic concentrations in NIF enclosures containing sol-coated optics," *NIF-0039225*.
3. Motyoyuki Suzuki, *Adsorption Engineering*, (Elsevier, New York, 1990), pp. 67–69.

6.5 Engineering High-Damage-Threshold NIF Polarizers and Mirrors

C. J. Stolz, F. Y. Genin, M. R. Kozlowski, Z. L. Wu

Polarizers and mirrors are required in the National Ignition Facility (NIF) to switch the pulse into and out of the main cavity and to direct the laser into the target chamber. Understanding the interaction of lasers with optical materials aids in the development of deposition techniques, coating designs, and post deposition processes necessary to manufacture optical coatings that survive the fluences required for ignition. The thin-film deposition technique affects the microstructure of coating layers, the formation of defects, and the interface strength between defects and the multilayer. The coating design determines the standing-wave electric field in the coating and the plasma durability of the outer surface of the coating. Post deposition processing by laser conditioning minimizes the effect of coating damage so that surface morphological changes remain stable with further irradiation. The effect of these damage sites on beam propagation must be understood to determine what size and type of damage is allowable.

6.5.1 Thin-Film Deposition Technology

A multitude of deposition techniques exist over a large range of deposition energies (0.01–5000 eV). The thin-film deposition technique must be carefully selected to meet the stringent coating requirements of the NIF. The optical substrates are large (up to 1 m diagonally) and numerous enough that the process should be scalable for coating multiple optics at a time. Spectral requirements are challenging, requiring complex thin-film designs with thickness control to less than 10 nm and full-aperture nonuniformities less than 1%. Coating stress must be minimized to prevent distortion of the laser beam wavefront. Finally, the coatings must not interact significantly with the laser during exposure to high fluences.

6.5.2 Electron-Beam Deposition

Electron-beam (e-beam) deposition, illustrated in Figure 6.5-1, is a technology that scales to large optics with sufficient precision to produce complex multilayer thin-film designs. This deposition technology has historically been the main tool of the laser-fusion community for making high reflectors and polarizers because laser-damage thresholds have historically increased with the requirements for higher ICF-laser energies. Significantly less effort has been put into alternate deposition processes, so they remain relatively unoptimized for high-fluence, large-aperture fusion optics.

The e-beam deposition process consists of an electron beam generated by thermionic emission from a filament. The beam is focused and directed by permanent magnets to a grounded source material. To minimize source-material drilling, electro-magnets are used to raster-scan the beam over the surface. A

reactive deposition environment provides greater flexibility in selecting source material compositions that interact more favorably with the e-beam.

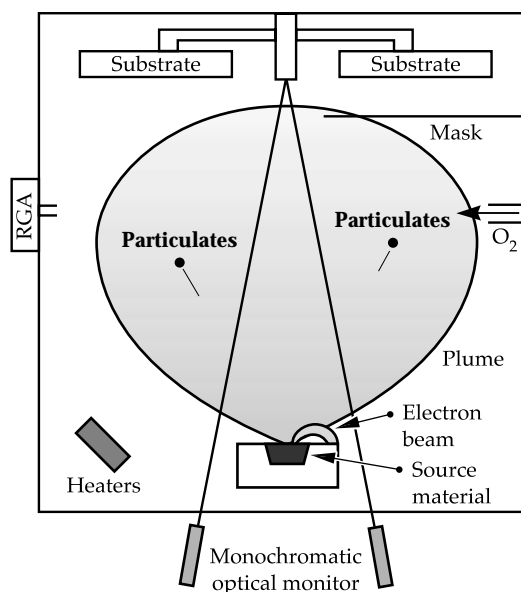


Figure 6.5-1. Typical e-beam deposition system.

The surface topography of the source material affects the stability and uniformity of the deposition plume (Figure 6.5-2) that has a cosine function distribution normal to the incident surface. An irregular surface leads to an unstable deposition plume that limits the layer-thickness control required to successfully coat the complex NIF designs. If the source material is an oxide, a small molten pool forms in the immediate vicinity of the e-beam due to the low thermal conductivity of the dielectric source material. The e-beam must be swept over the source material to prevent drilling or boring, which changes the deposition angle. Oxide source materials also tend to have an irregular surface because they are initially formed from multiple pellets. A metallic source material, deposited in a reactive environment to generate an oxide film, has a lower melting point and higher thermal conductivity. Therefore, a large, flat, molten metal surface with minimal drilling produces a stable deposition plume, as illustrated in Figure 6.5-2. A 3× improvement in plume stabilization is experimentally observed when switching from hafnia (HfO_2) to hafnium (Hf) source material.¹

The interaction of the e-beam with the starting material generates particles that can be incorporated into the film as coating defects.² Hafnia plugs have poor packing densities due to their low thermal conductivity. Small air voids in these plugs are exposed by the oscillating e-beam, resulting in ejected particles. Additionally, a phase transformation exists in hafnia at $\sim 1700^\circ\text{C}$. During evaporation the top surface of the

hafnia is molten at $\sim 2200^{\circ}\text{C}$ and the bottom and edge surfaces are in contact with a water-cooled crucible at $\sim 100^{\circ}\text{C}$. There is therefore a region in the material that undergoes a monoclinic to tetragonal phase transformation, resulting in a volumetric expansion of 3.8%. Volume-induced stresses, then, may also eject particles.




Figure 6.5-2. Source material composition affects interaction between the electron-beam and the source material, and the density of particle formation.

6.5.3 Ion-Beam Deposition Methods

In the quest to achieve thin films resistant to higher laser energies, other coating processes have been investigated for higher damage thresholds than the e-beam process. Two of the more promising methods that have been studied are ion-assisted deposition (IAD) and ion-beam sputtering (IBS). IAD uses a conventional e-beam source to generate the deposition plume and an ion gun pointed at the optical substrate to increase the energy of the deposited particles (Figure 6.5-3). IBS uses an ion source that strikes a target to sputter atoms with higher arrival energies onto the substrate (Figure 6.5-4). The higher energy of the deposited particles from these techniques generates denser films than e-beam coatings. These lower-porosity films are less sensitive to humidity-induced spectral shift and stress changes.³ Although environmental insensitivity is desirable from a spectral and stress perspective, the increased film stability causes defects to be more tightly bound resulting in catastrophic ejections during high-fluence irradiation. These damage sites generally grow upon continued irradiation, whereas most e-beam-coating nodular-ejection sites are stable at the NIF peak fluence.⁴




Figure 6.5-3. Typical IAD deposition system.




Figure 6.5-4. Typical IBS deposition system.

6.5.4 Laser Damage

Lasers create damage in optical thin films when they exceed the damage fluence of the coating. Figure 6.5-5 illustrates typical damage morphologies of hafnia-silica multilayer coatings. The four most common damage morphologies are:

- Delaminates.
- Flat-bottom pits.
- Nodular-ejection pits.
- Plasma scalds.

Only damage morphologies that grow, after continued laser exposure, to sufficient size to affect system performance must be prevented. Previous work has demonstrated that damage less than 280 μm in diameter will not degrade the performance of the laser.⁵ Because the components must survive multiple shots, we must understand the growth of the various damage morphologies to determine coating lifetimes (Figure 6.5-6).⁶ Therefore, careful characterization of these morphologies is necessary to determine which are fluence-limiting, how they affect the lifetime of an optical coating, and what are their sources.




Figure 6.5-5. SEM images of laser-induced damage morphologies in hafnia-silica multilayer mirror and polarizer coatings.

6.5.5 Delaminates

Delaminates grow at the lowest fluence and quickest rate, so improvements in the coating damage threshold must start with elimination of this morphology. Delaminates are caused by the removal of the outermost layer, most likely due to thermally induced mechanical stresses such as buckling. Several phenomena imply that a surface plasma causes the delamination damage morphology. Plasmas are always observed during irradiation in coating areas that delaminate, with a strong

correlation between plasma size and delaminate size. Delaminates can be found around coating defects and also in defect-free regions. The scanning-electron micrograph (SEM) in Figure 6.5-5 reveals signs of mechanical failure, such as curled and torn edges of the delaminated overlayer. The lack of plasma scalding in or around the delaminated site suggests that the scalded overlayer tears when the plasma extinguishes or when the plasma propagates away from the surface during overlayer ejection.

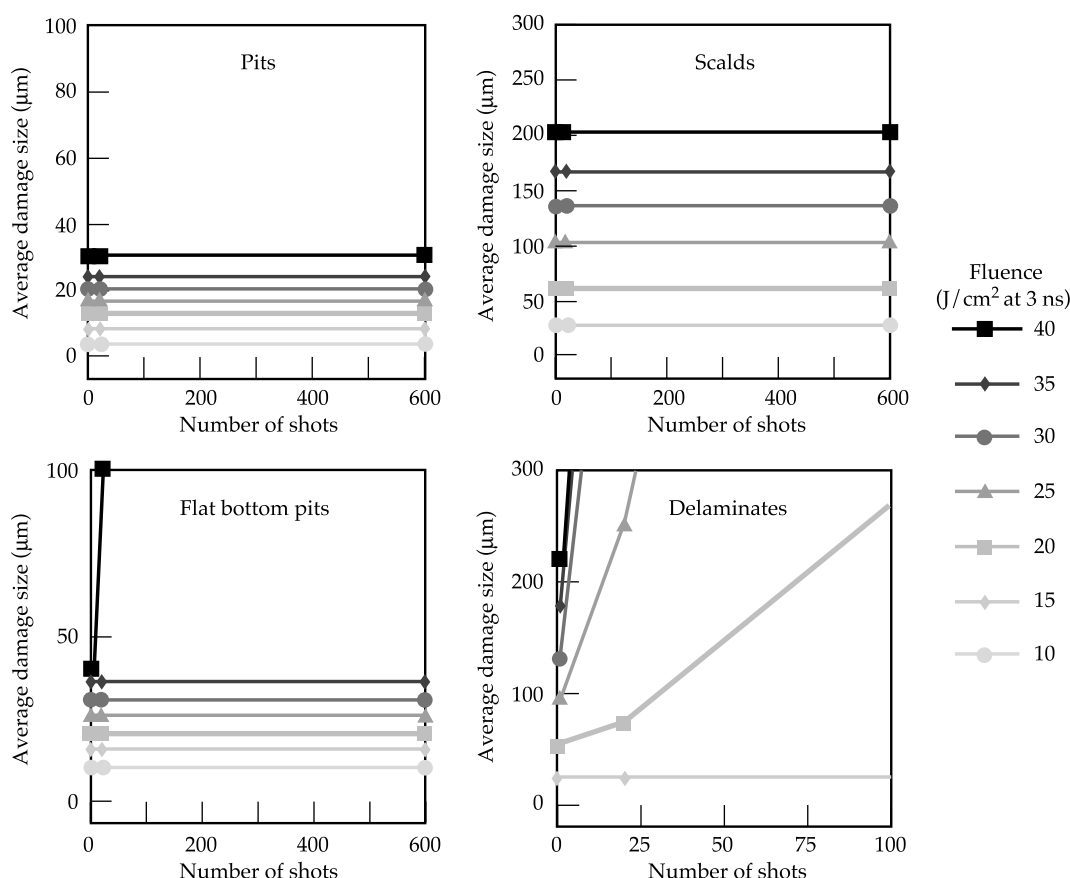


Figure 6.5-6. Growth of various damage morphologies exposed to multiple-shot irradiation for a range of fluences indicates that delaminates are unstable while the other morphologies are benign at NIF fluences of 10.5 J/cm² (polarizers) and 21.9 J/cm² (mirrors) at pulse lengths of 3 ns.

Delaminates in polarizers are affected by laser incident angle. Delaminates occur at Brewster's angle in the Beamlet-design coatings, yet they are eliminated at an incidence of 45 degrees, suggesting an electric-field dependence.

Maintaining a minimum overcoat thickness eliminates delaminates.⁷ A polarizer coating consists of three parts—a polarizing stack with antireflection coatings on both sides. The thickness of the outer low-index layer is determined by

optimization of the design for high transmission of “P” polarized light. Typical optimized edge-filter-design polarizer coatings, such as those used in the Beamlet polarizer, have overcoat optical thicknesses of about $\lambda/7$ (~150 nm physical thickness). These designs delaminate easily during laser exposure. Polarizer coatings with overcoats thicker than $\lambda/3$ (~300 nm physical thickness) do not show this delamination, even in the presence of plasmas from nodule ejections. Previous work has demonstrated that damage threshold improves for high reflectors with silica overcoats that have a $\lambda/2$ optical thickness.⁸ This type of overcoat does not affect the spectral characteristics of the coating at the design wavelength and therefore is an “absentee layer.” If the optimized design has a thin overlayer, an additional $\lambda/2$ can be added to the thickness to eliminate delaminates (as shown in Figure 6.5-7) without sacrificing optical performance.




Figure 6.5-7. Thicker polarizer overcoats significantly improve laser-damage morphology. In micrographs (a) and (b), delamination occurs within the scalds. In (c), no delamination occurs.

6.5.6 Flat-Bottom Pits

Flat-bottom pits are circular divots that fracture along interfaces correlating to electric-field maximums. They are interfacial in nature and typically occur within the first few outer layer pairs. Flat-bottom pits in hafnia-silica multilayer coatings grow at high fluences, so they could affect the optic lifetime. These pits are created around nodular-ejection sites as well as in the absence of visible defects. The current theory is that the presence of embedded nanometer-scale absorbing seeds is responsible for the formation of flat-bottom pits.⁹ This theory states that there is sufficient energy at electric-field peaks to create a plasma, leading to film buckling and localized radiation damage.^{10,11}

During development of hafnium deposition for high-damage-threshold coatings, it was found that flat-bottom pits did not occur in hafnium-deposited coatings irradiated at NIF fluences, but they did occur in hafnia-deposited coatings. Several factors could explain the dependence of the starting material composition on this damage morphology, including improvement of interfacial strength and elimination of nanometer-scale absorbing seeds. It is not surprising that the

interfacial quality in these coatings is material-dependent, because the hafnia layers are polycrystalline and rough while the silica layers are amorphous and smooth. However, starting material composition also affects interfacial quality; coatings deposited from hafnium and silica have fewer interfacial voids than those deposited from hafnia and silica (Figure 6.5-8). The effect of starting material composition on the presence of absorbing seeds is unknown.

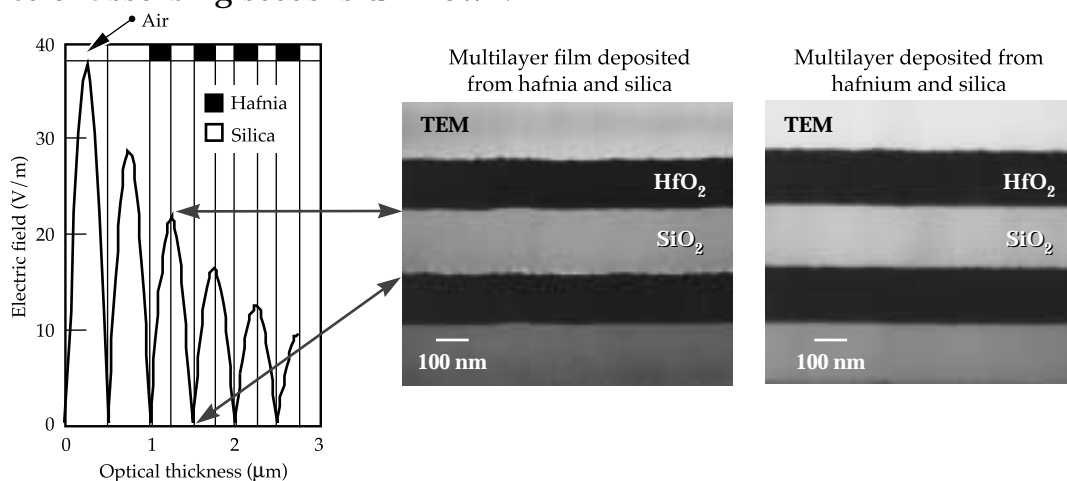


Figure 6.5-8. Interface quality as a function of starting material composition.

To better understand the influence of the two different interfaces, a modified high-reflector design was deposited to cause electric-field peaks at both types of interfaces.¹² Simultaneously, a standard high reflector was deposited as a control piece. The coatings were deposited using hafnium and silica starting materials. After damage testing, flat-bottom pits were not observed on the standard high reflector at up to twice the NIF fluence, although some interfacial damage occurred at the substrate/thin-film boundary. The modified design had flat-bottom pits due to the significantly higher electric fields, but only at the hafnium-over-silica interface (Figure 6.5-9).

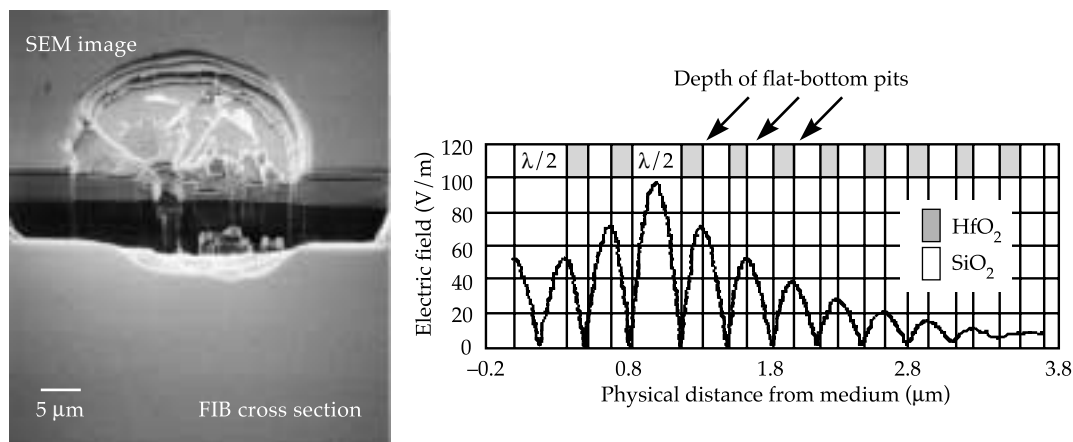


Figure 6.5-9. Flat-bottom pit depth correlates to hafnia grown over silica interfaces at electric field peaks.

6.5.7 Nodular-Ejection Pits

The ejection of nodular defects during laser exposure results in a crater-like pit. Modeling of nodular defects reveals up to a $5\times$ increase in the electric field due to the nodule geometry.¹³ These enhanced electric fields create localized, thermally induced stresses that are the likely cause of nodular ejection.

To understand the relationship between nodule size and damage threshold, an atomic force microscope (AFM) was used to characterize nodules before and after laser irradiation.¹⁴ A correlation was observed between nodule height (which is proportional to the defect-seed diameter) and damage threshold. The increase of electric-field enhancement with larger seed size correlates well with the observation that larger-diameter seeds have a lower damage threshold. This information can be used to optimize the coating processes so that defects are too small to damage the coatings at the operating fluence of the NIF laser.

The coating process is inherently unstable, resulting in the generation of particles that become embedded in the coating. The primary defect source in the coating process is material ejected from the source, but defects are also caused by material that flakes off shields around the chamber walls and the tooling, particles from the interaction of the e-beam with the source material, and particles from arcing of high-voltage surfaces.¹⁵ The e-beam coating process takes place at an elevated temperature around 200°C. The thermal expansion of stainless-steel tooling and aluminum shields is different than that of the coating materials, causing the coatings on the interior metal surfaces to delaminate. Fracturing and delamination can also be caused by intrinsic stresses in the coating material (due to the buildup of coatings over multiple runs) and weak adhesion of the coating material (due to the monolayers of adsorbed gasses and organic contaminants from the metal-cleaning process). All these particles can be transported by the thermal plume onto the optical surfaces being coated.¹⁶ Proper chamber preventive maintenance and cleanliness are required to produce high-damage-threshold coatings.

A focused ion beam (FIB) can be used to characterize the nodules to determine their origin.¹⁷ The FIB is positioned perpendicular to the nodule and raster-scanned across it; half of the nodule material is removed by sputtering to reveal a vertical cross section (Figure 6.5-10). Once half of the nodule is exposed, the FIB fluence can be decreased and positioned to an incidence angle of about 60 degrees from normal. By bombarding the cross-sectioned surface at this angle, secondary ions are generated and collected to produce high-resolution images of the nodule, which reveal information about the defect seed size, shape, and chemistry (Figure 6.5-11). This information can be used to infer the origin of the seed.

Round, smooth-edged seeds indicate molten ejecta from the interaction of the e-beam with the source material. Rough-edged seeds indicate solid ejecta from either e-beam/source interaction, arcing of coated surfaces, or particles that contaminated the substrate before coating. By comparing the contrast of the seed with the individual coating layers, particle composition can be inferred. Thick hafnia flakes

and thin multilayer flakes have also been observed, indicating that coating flakes from chamber surfaces can become incorporated into the coating during deposition.




Figure 6.5-10. Focused-ion-beam milling is used to cross-section coating defects for characterization of the seed.




Figure 6.5-11. Comparison of nodular defect seed geometries from various deposition techniques.

Defects in IBS coatings typically occur with a lower density than e-beam coatings. Laser-damage thresholds in small areas are typically very high ($>40 \text{ J/cm}^2$ at 3 ns), due to the low probability of irradiating a defect. Unfortunately, large-aperture damage thresholds with a higher probability of containing a fluence limiting defect are significantly lower ($\sim 10 \text{ J/cm}^2$ at 3 ns) than those of typical e-beam films of the same coating materials and optic size. One possible explanation is illustrated by the difference in the boundaries between the defects and the multilayer for e-beam and IBS coatings. IBS coatings typically have continuous boundaries, even for large defects, but large e-beam defects have significantly convoluted, weaker boundaries. Laser-interaction experiments on IBS coatings show that large defects ejected at low fluences grow considerably with repeated irradiation, but large ejected defects in e-beam coatings typically have significantly less damage that is more stable.

6.5.8 Laser Conditioning to Minimize Damage

The severity of pitting can be reduced by ejection of nodular defects at fluences lower than the peak NIF fluence. The process of irradiating the coating at lower fluence is termed laser conditioning. Figure 6.5-12 illustrates the laser conditioning method used to raster-scan a large-aperture optic, with each subsequent scan at a higher fluence. Nodular defects have different damage thresholds depending on the diameter and depth of the seed.¹⁸ The lower the laser fluence that is used to eject a nodular defect, the smaller the plasma that is created, reducing the probability of microcracks propagating from the pit and decreasing the size of the plasma scald. After the nodule is ejected, a high-damage-threshold pit remains because electric fields are low in smooth pits. Nodules ejected at higher fluences have a greater probability of generating microcracks, which on subsequent shots result in electric-field-enhancement sites for plasma formation that can lead to runaway growth.




Figure 6.5-12. Six-step laser conditioning process.

Large-aperture laser conditioning of coatings has been demonstrated to improve the damage threshold of optical coatings.¹⁹ To reduce the cost of laser conditioning, a single-step laser conditioning method was tested.²⁰ By exploiting the Gaussian

shape of the beam and raster scanning, the surface is exposed to a fluence ramp needed for conditioning (Figure 6.5-13). Beamlet mirrors were successfully conditioned with this technique.




Figure 6.5-13. Single-step laser conditioning process.

6.5.9 Plasma Scalds

A plasma scald is created when a plasma forms during irradiation and erodes the overlayer. In the eroded area, a rougher surface with higher surface energy is observed; the improved wettability of the scalded surface indicates surface cleaning or charging by the plasma. Previous work has shown that the damage threshold of a coating can be significantly increased by adding a $\lambda/2$ -thickness low-refractive-index layer, which does not affect the spectral performance of the coating.⁸ Although the reasons for the improvement in the damage threshold are not well understood, the difference in the magnitude of surface erosion due to the presence of a plasma during irradiation is significant. Upon repeated exposure to laser pulses, low-index material shows minimal erosion (~ 6 nm), but high-index material erodes completely through to the next layer (~ 140 nm; Figure 6.5-14).

The only observable, stable, damage morphology larger than the critical $280\text{-}\mu\text{m}$ size was plasma scalding, so the effect of the scalds on the laser beam was characterized. Plasma scalds are the result of partial erosion of the outer layer, so they can be modeled as micromirrors. To create sufficient focusing to overcome diffraction effects, the erosion thickness would have to penetrate the overlayer.²¹ Typical depths of plasma scalds with silica overcoats are less than 5% of a layer thickness, so they behave as scatter sites, not as focusing elements.

Measurements were also conducted to understand the influence of the increasing surface roughness on the reflectivity of the scalded area. A 50-mm -diameter plasma scald was clearly visible to the naked eye as a discoloration of the coating. However, reflectivity measurements at the NIF laser wavelength showed

no degradation within the resolution of the Nova photometer with an accuracy of $\pm 0.2\%$.




Figure 6.5-14. Plasma-scauld erosion depth is minimized with a silica overlayer.

6.5.10 Current Technology Status

Under the NIF coating-development program, laser-conditioned coatings typically have zero to four defects, over the Beamlet aperture, that limit the peak fluence. Previous laser-interaction studies were done on small samples because of their larger defect density, maintaining a statistically significant sampling of defects; these smaller samples are easily characterized under a variety of techniques, including AFM, SEM, and FIB. To characterize fluence-limiting defects, a nondestructive technique that can scan the entire surface will have to be used. Photothermal microscopy currently holds the greatest potential for nondestructively identifying fluence-limiting defects, as shown by the good correlation between high photothermal signal and low damage threshold illustrated in Figure 6.5-15.^{22,23} However, the current data-acquisition rate is so slow that scanning a NIF-size optic would take approximately 200 days.

If fluence-limiting defects are present on NIF optics, their effect on optic lifetime will be mitigated by several statistical factors. Because only a small percentage of the coated surface is exposed to the NIF peak fluence, damage growth will only occur when a low-density, fluence-limiting defect—which populates only $\sim 0.0005\%$ of the surface—is struck by the highest-fluence regions of the beam. The average growth rate of damage sites will be reduced because many of the NIF shots will be below

peak fluence. Finally, coatings will be conditioned at fluences slightly below rated peak fluences to minimize conditioning-induced damage, also slowing damage growth.




Figure 6.5-15. Photothermal signal correlates well with the laser-damage threshold.

6.5.11 Summary

Optical coatings can be generated by several energetic processes, including e-beam deposition, ion-assisted deposition, and ion-beam sputtering. E-beam coatings have demonstrated the highest large-aperture damage thresholds for fusion lasers. By understanding the sources of various damage morphologies, the coating process and design can be engineered to produce coatings capable of withstanding NIF fluences. Delamination can be eliminated by selection of an appropriate overcoat thickness. Flat-bottom pits can be eliminated through strengthening the layer interfaces by starting with hafnium instead of hafnia or reducing the peak electric fields at the layer interface. Nodules are reduced by good substrate-cleaning techniques, proper chamber maintenance, and the use of hafnium instead of hafnia to reduce ejected particles. The severity of nodular-ejection pits and plasma scalds can be minimized by laser conditioning—irradiating at successively higher fluences to induce nodular ejection at fluences below operational values—to reduce plasma formation and microcracking.

Notes and References

1. C. J. Stolz, M. K. Von Gunten, R. P. Bevis, and D. Smith, "The advantages of evaporation of hafnium in a reactive environment to manufacture high damage threshold multilayer coatings by electron beam deposition," in *Proc. SPIE*, to be published.
2. R. Chow et al., *Appl. Opt.* **32**, 5567–5574 (1993).
3. P. J. Martin et al., *Appl. Opt.* **22**, 178–184 (1983).
4. C. J. Stolz et al., *Proc. SPIE* **2714**, 351–358 (1996).
5. J. T. Hunt, K. R. Manes, and P. A. Renard, *Appl. Opt.* **32**, 5973–5982 (1993).
6. F. Y. Genin, C. J. Stolz, and M. R. Kozlowski, *Proc. SPIE* **2966**, 273–282 (1996).
7. C. J. Stolz et al., *Proc. SPIE* **2966**, 265–272 (1996).
8. B. E. Newnam, "Laser-induced-damage phenomena in dielectric films, solids and inorganic liquids," Ph.D. dissertation, Univ. of Southern California (1973).
9. J. Dijon et al., *Proc. SPIE* **2714**, 416–424 (1996).
10. J. Dijon, T. Poiroux, and C. Desrumaux, *Proc. SPIE* **2966**, 315–325 (1997).
11. J. Dijon, G. Ravel, and B. André, *Proc. SPIE*, **3578**, 398–407 (1998).
12. S. C. Weakley et al., *Proc. SPIE*, **3578**, 137–143 (1998).
13. R. H. Sawicki, C. C. Shang, and T. L. Swatowski, *Proc. SPIE* **2428**, 333–342 (1995).
14. R. J. Tench, R. Chow, and M. R. Kozlowski, *J. Vac. Sci. Technol. A* **12**, 2808–2813 (1994).
15. F. L. Williams, G. A. Peterson, R. A. Schmell, and C. K. Carniglia, *Proc. SPIE* **1624**, 256–270 (1991).
16. S. Matteson and R. A. Bowling, *J. Vac. Sci. Technol. A* **6**, 2504–2507 (1988).
17. C. J. Stolz, R. J. Tench, M. R. Kozlowski, and A. Fornier, *Proc. SPIE* **2714**, 374–382 (1995).
18. M. R. Kozlowski, R. J. Tench, R. Chow, and L. M. Sheehan, *Proc. SPIE* **2253**, 743–750 (1994).
19. L. M. Sheehan, M. R. Kozlowski, F. Rainer, and M. C. Staggs, *Proc. SPIE* **2114**, 559–568 (1994).
20. C. J. Stolz et al., *Proc. SPIE*, **3578**, 144–153 (1998).
21. F. Y. Genin and C. J. Stolz, *Proc. SPIE* **2870**, 439–448 (1996).
22. C. J. Stolz et al., *Proc. SPIE* **3244**, 475–483 (1998).
23. Z. L. Wu et al., "Characterization and laser damage testing of fluence limiting defects in hafnia/silica multilayer coatings developed for the National Ignition Facility," to be published.

6.6 Diffractive Optics for the NIF

J. A. Britten, S. N. Dixit, S. M. Herman, M. C. Rushford, L. J. Summers, T. G. Parham, L. Auyang, I. M. Barton, B. W. Shore

The National Ignition Facility's (NIF's) baseline laser design incorporates three diffractive structures in the third-harmonic (351-nm) final optics assembly, as shown in Figure 6.6-1. These are (1) a focusing beam sampling grating (BSG)¹ for sending a known fraction (0.2 to 0.4%) of the transmitted light into a calorimeter for energy diagnostics; (2) a color separation grating (CSG)^{2–5} that transmits, with high-efficiency, third-harmonic light to the target, while directing the unconverted first and second harmonic light away from it, and; (3) a kinoform phase plate (KPP)⁶ that generates a tailored focal spot shape at the target plane. The optics' apertures are nominally 40 cm square. To withstand the design fluence, the diffractive structures' patterns must be etched into the bulk fused-silica substrate. Wet chemical etching by a buffered hydrofluoric acid (HF) solution was chosen as the pattern transfer method because (1) the low aspect ratio of these structures is conducive to this process; (2) the process is inherently spatially uniform due to the kinetic control of the dissolution,^{7,8} and; (3) it is inexpensive to implement. A flowchart of the processing steps used to manufacture the optics is shown in Figure 6.6-2.




Figure 6.6-1. Schematic diagram of a NIF final optics assembly, showing the diffractive optics' location, element geometry, and specifications.

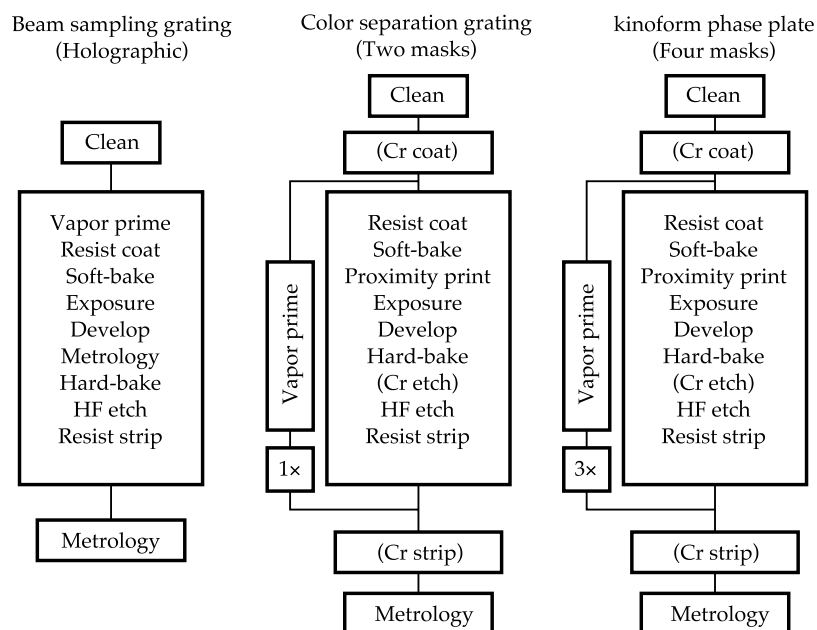


Figure 6.6-2. Flow diagram of the processing steps required to fabricate the three NIF 3 ω diffractive optics.

The BSG is a lamellar grating with a nominal period of from 1 to 3 μm and a modulation depth of about 20 nm. It is made holographically by projecting two interfering coherent spherical waves onto the photoresist-coated substrate in the appropriate geometry. (These two waves simulate the main beam being focused onto the target and the sampled beam being focused into a calorimeter.) The resulting latent image is developed to give a grating mask in resist. This pattern is transferred to the surface of the substrate by etching the exposed areas between the resist grating lines with HF solution.

The CSG is a staircase grating design.^{2–5} It is made by proximity printing one line of the period through a chrome-on-quartz master mask onto a photoresist layer on the target substrate, developing this pattern, and transfer etching it (again, with HF solution) to a precise depth (equal to one wave of optical phase difference in transmission at the chosen use angle and wavelength). This process is repeated to generate the second line of the period—using the same mask (or a different one) offset laterally, relative to the first line etched.

The KPP is similarly made, by using a four-mask process and transfer wet etching⁶ to produce 16 step levels that approximate a continuous, irregular topography over a scale length of several millimeters, with a maximum optical path difference of one wave.

To manufacture diffractive optics plates (DOPs) onsite, we constructed a 2400-ft² facility at Lawrence Livermore National Laboratory (LLNL), which became

functional in January 1998. During 1998 we fabricated all three types of DOP at NIF sizes, using custom-designed prototype processing equipment and the manufacturing techniques that will be employed for full-scale production. This article focuses on the manufacture and performance of the beam sampling and color separation gratings, two optics that were not previously fabricated at NIF scales. We have made (and reported on) several KPPs for LLNL's Nova laser and other laser systems⁶ with processing methods that are identical to those used for the CSG. We also discuss the effects of standard antireflective (AR) coatings on the performance and laser damage characteristics of the diffractive optics, and describe efforts to combine the functionalities of several diffractive elements onto one surface.

6.6.1 The Beam Sampling Grating

A laser interference lithography facility has been constructed at LLNL to fabricate NIF BSGs. Figure 6.6-3 shows the geometry of the exposure system. A CW laser with a very long coherence length is required to generate stable, high-contrast fringes in the target plane. We use a 351-nm, Ar-ion laser because its wavelength very closely matches the third harmonic of the NIF. Thus, we can set up exposure geometries that are identical to the deployment geometries in the NIF's final optics assembly and obtain a high-quality focal spot of the diffracted beam with minimal chromatic aberration of the diffracted focus. The laser output in single-mode TEM_{00} operation is about 1.2 W. The exposure table is totally enclosed and is isolated from vibrations by pneumatic supports. Active fringe stabilization is used to ensure high pattern contrast. The major difficulty with this setup is finding suitable high-magnification objectives for transmitting the 351-nm light for the fast-focus sampling beam. All of the multielement objectives tested to date imprint some degree of modulation onto the fringe pattern from defects in the multielement structure, which are difficult to clean and eliminate by spatial filtering.

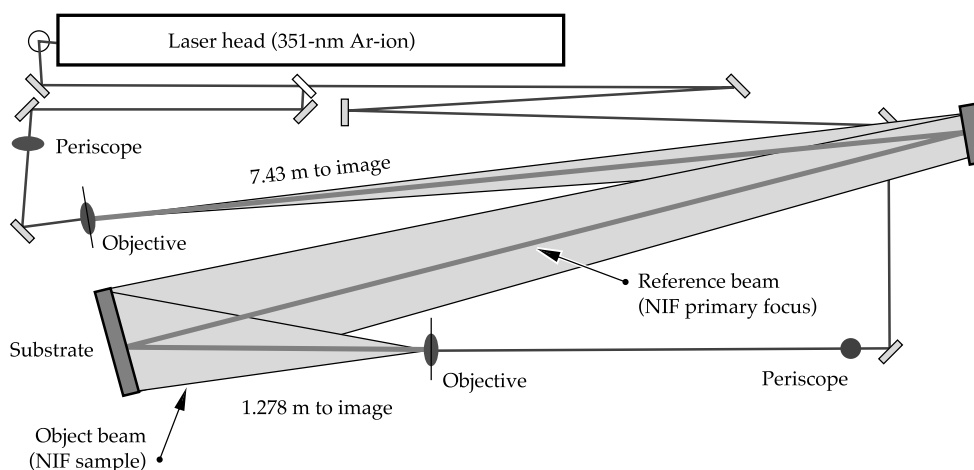


Figure 6.6-3. Schematic of a NIF beam sampling grating's holographic exposure geometry.

Using the process outlined in Figure 6.6-2, we have made demonstration BSGs on 41- by 39-cm fused-silica plates. A map of the first-order (focused) transmission efficiency of one such grating is shown in Figure 6.6-4. This measurement was made at 351 nm with the optic rotated 15° about the vertical axis with respect to the incident beam and with the diffracted beam at an angle of 15° above the optic's horizontal axis; the mean efficiency of 0.37% is within specifications. The spatial uniformity needs to be improved, even though model simulations giving the projected uncertainty of the spatial intensity variations for a NIF beamline suggest that this level of nonuniformity will not contribute substantially to the measurement error. The efficiency variation results from a nonuniform grating linewidth, not an etch-depth variation. The modulation seen in the diffraction efficiency is directly attributable to intensity variations in the spatially filtered, sample-writing beam, using a $40\times$ objective. Improvements to this uniformity will be possible with higher quality objectives.




Figure 6.6-4. Diffraction efficiency (%) at -1 transmitted order, for 351 nm at a 14° incidence angle. The beam sampling grating was manufactured on a 41- \times 39-cm fused-silica substrate. The mean efficiency is 0.37%, with a standard deviation of 0.07%.

6.6.2 The Color Separation Grating

Several CSGs were fabricated and tested for optical performance. Made by LLNL's laser plotter, the grating mask used for this process consisted of alternating $113\text{-}\mu\text{m}$ -wide clear lines separated by $232\text{-}\mu\text{m}$ -wide chrome spaces, 370-mm long, repeated over a 380-mm field. This pattern was printed on a 65-cm -diam fused-silica substrate. The mask was printed pixel-by-pixel by positioning a 113- by $500\text{-}\mu\text{m}$ -wide aperture with sub- μm accuracy and then, by opening a shutter, exposing the mask (previously coated with an evaporated chrome layer and a photoresist layer) to light from a 413-nm Kr-ion laser. The pattern was transferred to the mask by subsequently developing the resist and etching the exposed chrome.

The CSG target substrates were first coated with chrome and photoresist, and were then exposed through the grating mask to the output of a collimated Hg lamp. Next, the substrates were developed, hard-baked, and chrome-etched. Then, to make

the first step, they were HF-etched to a depth of 715 ± 10 nm. To etch the second step (while doubling the depth of the first step) the entire process was repeated, after the mask was laterally shifted 114 ± 1 μm (with the aid of fiducial marks on the mask and by viewing the substrate through confocal microscopes on the mask aligner). During this process, the one- to two- μm undercutting (edge erosion) of the chrome and fused silica edges was taken into account by making the open area of the master mask undersized. The end result was a grating, composed of a series of repeating stairsteps with a 345- μm period, across the clear aperture of the 41- by 39-cm substrate.

A map of the zero-order transmission efficiency of one CSG at 351 nm and 13° incidence angle is shown in Figure 6.6-5. The average efficiency was 90.5%. This optic had no antireflective coating, so the maximum theoretical efficiency, accounting for Fresnel losses, is 92.5%. The additional losses, about 2%, are distributed in many higher transmitted diffraction orders. The zero-order transmissions of this optic at 527 nm and 1053 nm are 0.05% and 0.26%, respectively, which exceed NIF specifications for 1ω and 2ω light rejection. We have made several such gratings with similar results.




Figure 6.6-5. Zero-order transmission efficiency at 351 nm for a 41- by 39-cm color separation grating (#O2Z) at a 13° incidence angle. The mean efficiency is 90.5% with a $\sigma = 0.75\%$. The zero-order transmission efficiencies at 1ω and 2ω were 0.26 and 0.05%, respectively. Both surfaces were uncoated.

6.6.3 The Effects of Sol-Gel Antireflective Coatings

Color Separation Gratings. Transmissive optics for LLNL's high-power lasers have for years been coated with high-damage-resistant sol-gel, colloidal-silica, antireflective (AR) films⁹ applied by dip coating. Table 6.6-1 gives the zero-order 351-nm diffraction efficiency of a NIF-sized CSG, which had its AR coating removed in places on the grating and on the back sides by being gently wiped with a wet soft cloth. The highest efficiency is obtained with the AR coating removed only from the grating surface. When the AR coating is on the grating surface, the reduction in back reflection is more than compensated for by an increased efficiency for the higher transmitted orders. Thus, the zero-order transmission efficiency drops.

Table 6.6-1. Zero-order 351-nm diffraction efficiency of a NIF-sized CSG.

	Both sides bare	AR-coated back side	AR-coated both sides
Zero-order transmission at 351 nm (%)	90.8	93.5	89.6

This drop in transmitted zero-order efficiency upon being AR coated results from a partial planarization of the etched structures, caused by surface tension effects as the coating is drying. Recall that the zero-order transmission at 351 nm is maximized when the etched structures are precisely deep enough to have integral wave multiples of phase retardation between adjacent elements. Planarization of the AR coating significantly alters this relation between adjacent steps. Figure 6.6-6 compares scanning electron micrographs (SEMs) of bare and AR-coated CSG steps of 1.4 and 0.7 μm (two-wave and one-wave steps). The ideal edge for a wet-etched CSG step is a quarter circle with a radius equal to the etched depth. The bare CSG steps that are shown closely approximate this ideal. The accumulation of colloidal sol-gel particles at the step edges (caused by capillary forces that pull liquid from adjacent areas while the coating is drying but still fluid) results in an increased film thickness many microns away from the edges. In the immediate vicinity of the steps, the coating is sufficiently thick to craze as it dries, due to shrinkage-related tensile stresses.




Figure 6.6-6. Scanning electron micrographs of (a) two-wave and (b) one-wave CSG steps wet etched into bare fused silica (upper) and into fused silica with a sol-gel AR dip coat (lower).

As a consequence of this effect, sol-gel AR-coated CSGs have significantly worse laser damage characteristics than when left uncoated. Modulation enhancement,

due to increased energy in high transmitted orders, causes damage to downstream optics under conditions where minimal damage would normally be expected. Because of the extreme coarseness of the grating period with respect to the wavelength and the size of the beam, this modulation exists for several meters downstream where the orders overlap. Figure 6.6-7 shows a SEM of output surface damage caused by the 351-nm, 7-ns irradiation of a NIF-sized, AR dip-coated optic with a CSG pattern on the input surface. The damage manifests itself as pinpoints that coalesce into lines that correlate with the grating period.




Figure 6.6-7. A micrograph of output surface damage on fused silica, caused by modulation from an AR dip-coated color-separation-grating pattern on an input surface that was illuminated with 351-nm light at 18 J/cm² from a 1.1-mm rastered beam at 7.5 ns.

Calculations have been made, using surface profiles of bare and coated CSG steps (as measured by atomic-force and scanning-electron microscopy), to estimate the modulation that imperfect step edges and AR coatings cause. These calculations consider infinitely periodic structures with a resolution of about 0.2 μm, and so capture both the macroscopic effects of the grating and the contributions of single grating steps. However, they do not include the difference in refractive index between substrate and coating. Figure 6.6-8 shows the predicted peak-to-average intensity ratio as a function of distance downstream from the surface of various CSGs. A NIF-sized CSG with a period of 345 μm and dip-coated AR coatings is large enough so that maximum intensity modulations of about 2.5 are predicted for several tens of millimeters downstream from it. The bare grating itself exhibits a modulation enhancement of about 1.5, entirely due to the nonvertical edges of the wet-etched steps. (The calculations assumed ideal etched depths. Errors in the depths of the etched steps are another source of intensity enhancement.) The modulation can be decreased by decreasing the density of the edges (that is, by increasing the period of the CSG). This model predicts that an uncoated CSG with an 1100-μm period will cause an intensity enhancement of about 1.3.




Figure 6.6-8. The calculated maximum peak-to-average intensity ratio, as a function of distance downstream from a wet-etched color separation grating (with and without a dip-coated AR layer). The calculations were done using measured profiles for the step edges and AR overcoat topography. The calculations did not take into account the refractive index difference between the coating and substrate. (a) = 345- μm grating period with an AR dip coat; (b) = 345 μm grating period, bare; (c) = 1100- μm grating period, bare.

Color separation gratings with vertical edges and perfectly etched step depths would not contribute to intensity enhancement. Ion-beam etching can potentially create substantially vertical sidewalls; but, as this is written, ion-beam etchers that can maintain an etch-depth uniformity of 1% over a 40-cm square aperture on a thick fused-silica substrate are not a proven technology.

As previously stated, dip coating is the baseline process for applying AR coatings to NIF transmissive optics. However, other available coating technologies have been investigated. One, known as spin coating, is a standard coating process in the semiconductor industry using high centrifugation to apply uniform films onto wafer surfaces. It has applied AR coatings that are significantly more conformal (less planarizing) than dip coating on small (5-cm-diam) CSGs. However, spin-coated AR layers on large rectangular parts, relevant to NIF, have exhibited significant and highly variable planarization effects, which depend on the spin rate, orientation of the grating steps, and their position on the substrate. It has been concluded that sol-gel coatings deposited by either method on the CSG's surface are unsuitable for NIF because of the likelihood of damaging downstream optics. An AR coating applied by vacuum-deposition techniques would be largely conformal and would not contribute to downstream modulation; but such coatings that can survive the NIF's 3ω baseline fluence do not exist.

The design for the NIF CSG is evolving toward using the natural dispersion of the focus lens to eliminate 1ω and 2ω light from the target, for all but the central portion of the beam. Here, a subaperture, uncoated CSG with a maximum allowable period (about 1 mm) will divert the 1ω and 2ω light from the target. A schematic of this concept is shown in Figure 6.6-9. This design will minimize the area and density of etched steps and will therefore minimize the modulation. The grating period is maximized by using a “split” grating design with the grating’s stairsteps facing away from the substrate’s centerline, instead of facing in the same direction across the entire aperture. With this design, light of a particular wavelength does not have to cross from one side of the grating to the other at focus, so the deflection angle needed becomes less, and the period can therefore become larger.




Figure 6.6-9. Schematic of a subaperture, split, color separation grating design that minimizes the area of the CSG and the etched step density.

Beam Sampling Gratings. The performance of a BSG is also influenced by AR coatings. The measured first-order diffraction efficiency of one NIF-scale BSG changed, from 0.4% bare to 0.15%, after a dip-coated AR layer was applied. The BSG structures are significantly shallower than the nominal 3ω AR-coating thickness of 72 nm. Calculations predict that transmitted first-order diffraction efficiency drops by more than a factor of three when the grating is covered with a completely planar AR overcoat. However, it does not change if the coating is completely conformal. These situations are illustrated in Figure 6.6-10. These data imply that AR dip coats on the BSGs are indeed largely planar.




Figure 6.6-10. Illustration of the calculated effect of sol-gel coating conformality on the first-order diffraction efficiency for a lamellar grating (representative of the NIF baseline design for the beam separation grating). At one extreme of conformality, there is no effect on the grating's efficiency, while at the other extremem (complete planarization) the efficiency is reduced by more than a factor of three.

The diffraction efficiency of a BSG drops about 40% when the same sol-gel layer is applied by spin coating at about 2000 rpm, suggesting an intermediate degree of planarization. In the case of BSGs, this planarization does not increase the likelihood of laser-induced damage, but the sensitivity of the diffraction efficiency to the coating's characteristics raises the possibility that environmental changes to the coating can cause unacceptable changes to the diffraction efficiency. Work is under way to characterize, both experimentally and via modeling, the sensitivity of the diffraction efficiency of AR-coated BSGs to environmental changes (such as those resulting from the presence of humidity and condensable organic contaminants).

6.6.4 Combining Diffractive Structures on One Surface

We have demonstrated the feasibility of combining the BSG and the CSG on the same surface¹⁰ by the wet-etch process. The fine BSG structures are formed first on a featureless substrate, then the CSG is made. (Fabricating the CSG first and then adding the BSG is problematic. Planarization of the resist over the CSG's topography would result in resist film-thickness variations beyond the processing latitude of the holographic exposure technique used to write the BSG pattern.)

The processing steps for both elements of the combined structure are the same as described earlier. Figure 6.6-11 shows CSG features that have been profiled by white light interferometry and BSG features that were measured by atomic force microscopy. The measured optical performance of each diffractive structure is not affected by the presence of the other. The slight tapering of the BSG profiles as they are propagated into the bulk during the CSG step etching is predicted by isotropic etching models.^{10,11} These model calculations show that the nominal NIF BSG

profile can be wet-etched approximately four- μm deep before the grating ridges begin to lose height. Since the combined depth of a CSG/KPP is only 2.3 μm at most and the methods for fabricating the KPP and CSG are the same, it would be straightforward to fabricate all three NIF diffractive structures on the same surface. This is an enabling development, since it might become important to concentrate all diffractive structures onto one surface and leave it bare, in light of the problems that have arisen in finding a suitable AR coating.

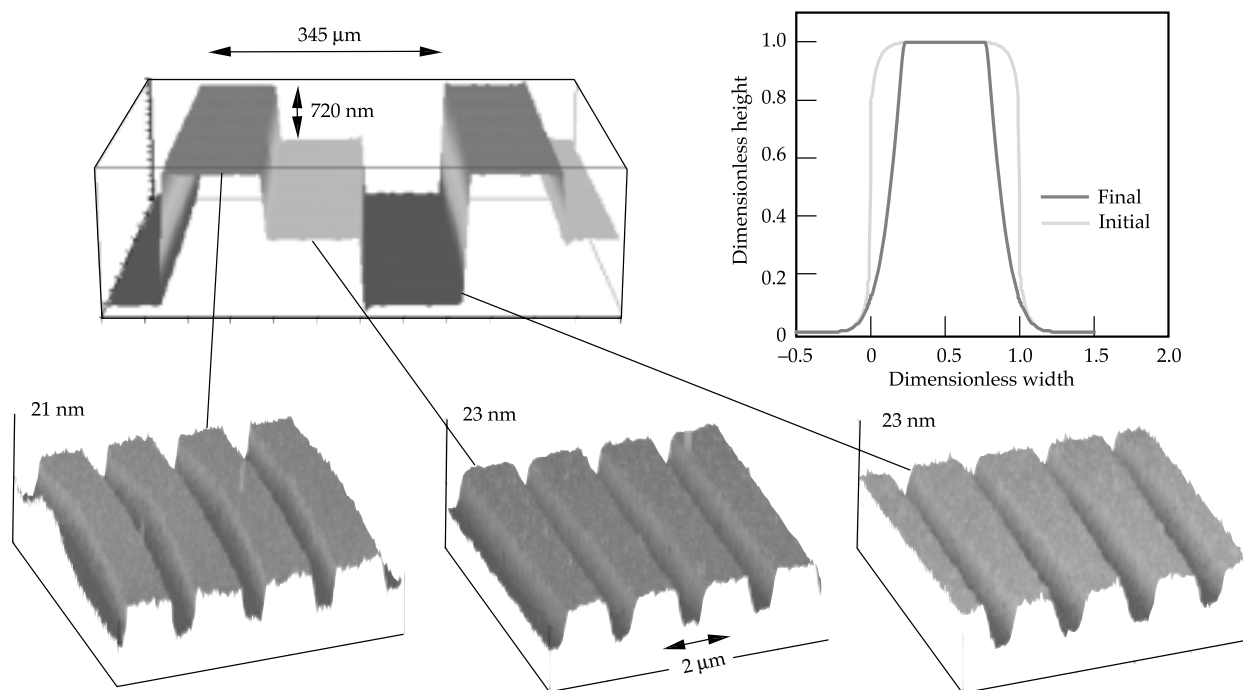


Figure 6.6-11. Beam sampling grating profiles measured by atomic force microscopy on different CSG steps, which are imaged by white light interferometry. Initial grating aspect ratio (height/width) ≈ 0.015 . Upper-right corner shows model calculations for the initial and final BSG shape after etching to the depth of the deepest CSG structure.

6.6.5 Summary

We have fabricated demonstration diffractive optics required for the NIF's baseline design at full-scale via a wet-chemical etching of fused silica, using fabrication techniques and processing equipment suitable for large-scale production. We have examined the effect of AR sol-gel coatings on the performance and laser-damage resistance of diffractive optics, and have concluded that the color separation grating must be left bare or the grating must be stripped of its sol-gel coating to minimize modulation effects that will damage downstream optics. We have also demonstrated the feasibility of combining all NIF diffractive structures onto one surface.

Notes and References

1. J. A. Britten, R. D. Boyd, M. D. Perry, B. W. Shore and I. M. Thomas, "Low-Efficiency Gratings for Third-Harmonic Diagnostic Applications," 1st Int'l. Conf. on Solid State Lasers for Application to Inertial Confinement Fusion, M. Andre and H. T. Powell, eds., *Proc. SPIE* **2633**, 121–128 (1995).
2. S. N. Dixit et al., "Color Separation Gratings for Diverting the Unconverted Light Away from the NIF Target," 2nd Int'l. Conf. on Solid State Lasers for Application to Inertial Confinement Fusion, M. Andre, ed., *Proc. SPIE* **3047**, 463–470 (1996).
3. H. Dammann, *Appl. Opt.* **17**, 2273–2279 (1978).
4. M. W. Farn, M. B. Stern, W. B. Veldkamp, and S. S. Medeiros, *Optics Lett.* **18**, 1214–1216 (1993).
5. T. H. Bett et al., "Diffractive Optics Development for Application on High-Power Solid State Lasers," 1st Int'l. Conf. on Solid State Lasers for Application to Inertial Confinement Fusion, M. Andre and H. T. Powell, eds., *Proc. SPIE* **2633**, 129–140 (1995).
6. M. C. Rushford, S. N. Dixit, I. M. Thomas, A. M. Martin, and M. D. Perry, "Large-Aperture Kinoform Phase Plates in Fused Silica for Spatial Beam Smoothing on Nova and the Beamlet Lasers," 2nd Int'l. Conf. on Solid State Lasers for Application to Inertial Confinement Fusion, M. Andre, ed., *Proc. SPIE* **3047**, 282–292 (1996).
7. G. A. C. M. Spierings, *J. Mater. Sci.* **28** 6261–6273 (1993).
8. D. T. Liang and D. W. Readey, *J. Am Ceram. Soc.* **70**, 570–577 (1987).
9. I. M. Thomas, *Appl. Opt.* **25**, 1481–1483 (1986).
10. J. A. Britten and L. J. Summers, *Appl. Opt.* **37**, 7049–7054 (1998).
11. S. Lindau, *Opt. Acta* **29**, 1371–1381 (1982).

7.0 THE MISSION-FIRST DEPLOYMENT STRATEGY FOR THE NIF..... 1

7.0 The Mission-First Deployment Strategy for the NIF

Brian MacGowan

In March 1999, members of NIF Project together with representatives from NIF user groups began to examine the beam installation plan for the NIF to determine if it could be optimized to better meet users' mission needs throughout the deployment period. The recommendations that resulted from those investigations have become known as the "Mission-First" deployment strategy, a strategy that incorporates flexibility in the sequence in which beams are commissioned and become available for experiments. This flexibility allows facility capabilities such as symmetric beam geometries to be available when determined by mission priorities rather than by construction schedules.

The laser hardware for the NIF is located in two laser bays, each of which provides beams to one hemisphere of the target chamber. Prior to 1999, the deployment strategy for the NIF was based on installation of laser beam infrastructure and laser line-replaceable units (LRUs) sequentially across each laser bay with simultaneous operation and construction. This strategy was described in the ICF/NIF Transition Plan (ref: LLNL's Transition Plan (7/98) Nova Operations to NIF Operations 1999–2004, 7/23/98). A major disadvantage with this strategy was the delaying of symmetric experiments, requiring beams from both laser bays, until installation of all beamlines. The early first-bundle strategy that installed temporary infrastructure for one bundle of 8 beams in Laser Bay 2 delayed eventual installation of the other beamlines in that Laser Bay (see Figure 7-1). The safety aspects of operating and constructing beam infrastructure at the same time also raised concerns.




Figure 7-1. The baseline deployment sequence for NIF beams reproduced from the ICF/NIF Transition Plan.

The Mission-First strategy relies on eliminating the temporary first bundle in favor of completing the beam path infrastructure for Laser Bay 2 as soon as possible, and that for Laser Bay 1 soon after. The beamline infrastructure includes amplifier buses, cable plant, gas-handling utilities, spatial filter vacuum vessels, beam tubes, mirror support structures, and Final Optics Assemblies. Once this infrastructure has been installed and commissioned, the contractors responsible can then vacate the facility, and installation of the laser components can begin. To commission beamlines requires installation, alignment, and testing of beam components that are designed as LRUs. The maintenance philosophy for the NIF is to remove and install LRUs during normal shot operations. By completing the infrastructure prior to LRU installation, the safety issues of concurrent infrastructure construction, LRU installation, and shot operations can be avoided by combining only LRU installation and shot operations. Apart from the volume of LRU installation, there should be no basic difference between activities in the NIF during the installation phase and during long-term operations.

By completing all infrastructure prior to beam commissioning, the Mission-First strategy allows the sequence of beam installation to be optimized to meet users' mission needs. The infrastructure can be considered as a chandelier into which the light bulbs, the LRUs, can be installed in any sequence. Experiments during the deployment phase of the NIF can be broken down into three broad categories summarized in Figure 7-2.

1. Experiments that require symmetric illumination of a hohlraum oriented such that beams enter through upper and lower laser entrance holes (LEHs) and have a rotational symmetry about the vertical. Typical experiments use the symmetric x-ray field in the hohlraum to implode spherical capsules (see Figure 7-2a)
2. Experiments that require less symmetry but a large amount of energy in beams incident through one LEH of a half hohlraum or "halfraum." Typical experiments use the x-radiation in the halfraum to drive shocks and other phenomena in a package at the far end of the halfraum (see Figure 7-2b). These experiments can be performed with the hohlraum axis either horizontal or vertical.
3. Experiments that use the beams to directly illuminate targets and generate pressure from material ablation in order to drive shocks and accelerate thin foils (see Figure 7-2c). These experiments benefit from beams that are grouped together in a small solid angle close to the foil normal.

For each of these experimental geometries other beams will be used to illuminate separate x-ray emitting plasmas (backlighters) that are used to radiograph the main experiment. (The diagrams show colored "quads" incident on various targets. Each colored quad represents 4 NIF beams, quads of the same color have 4-fold rotational symmetry about the vertical).




Figure 7-2. Three of the target geometries to be employed by NIF users, (a) vertical hohlraum with azimuthally symmetric beams from the top and bottom of the target chamber, (b) half hohlraum (halfraum) with a single cluster of beams incident on the LEH from the top and bottom of the target chamber or with symmetric beams incident from the bottom of the target chamber, and (c) planar foil directly illuminated by beams from the bottom of the target chamber. (Each “quad” in the figure represents four NIF beams).

The application of these different target geometries to different user campaigns has been discussed and described at length in many reviews of the Stockpile Stewardship campaigns (e.g., NIF Program Review Committee, 4/3/00–4/4/00). The important point to make here is that the Mission-First strategy, with early installation of infrastructure, allows great flexibility in choosing a deployment sequence that satisfies the maximum number of users at any point in time. Mission-First retains that flexibility throughout the deployment so that, as mission priorities change, so can the facility capabilities.

Figure 7-3 summarizes the current deployment sequence that has been discussed extensively with user groups over the past year, (for example, at the Campaign 4, 7, and 11 Program Manager’s meeting 2/8/00). This sequence has not been optimized and will probably change as the NIF rebaseline and associated schedule firms up. NIF Mission Support has a milestone to work with user groups to develop a consensus for the initial deployment sequence by January 2001.

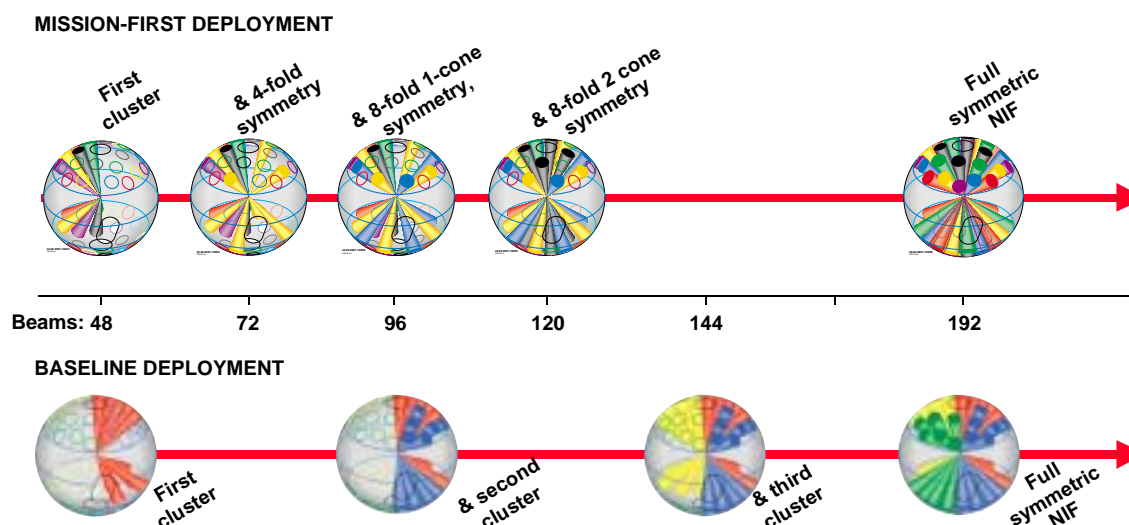


Figure 7-3. Mission-First and baseline deployment sequence, quads of 4 NIF beams are shown as colored cones within the NIF target chamber.

The provisional Mission-First sequence as shown at the top of Figure 7-3 begins with the commissioning of a single cluster of 48 beams. This configuration will allow experiments with horizontal hohlraums to begin with up to 200 kJ of energy in the hohlraum in 10 quads and up to 40 kJ available in backlighter beams. The early installation of a full cluster allows weapons physics experiments to start in this horizontal halfraum geometry.

With completion of full infrastructure in both laser bays, we can then add beams in groups with 4-fold rotational symmetry about the vertical. The sequence, as shown in Figure 7-3, adds 8 quads of beams that are all at 44.5° to the vertical (shown in yellow in Figure 7-3). This configuration allows experiments that require symmetry to begin in vertical hohlraums with up to 160 kJ of energy in the 32 symmetric beams, with the remaining beams from the first cluster available as backlighters. Early ignition experiments to develop symmetry measurement techniques can begin as well as weapons physics experiments that require symmetry, such as convergent mix and high-temperature hohlraum development experiments.

The next step is to add a further set of 4-fold symmetric, 44.5° beams (shown in blue in Figure 7-3) providing an 8-fold rotational symmetry with up to 320 kJ available for vertical hohlraum experiments with 64 symmetric beams and the other 32 beams available for backlighting. The increasing symmetric energy allows all experiments requiring symmetry to move to more interesting regimes.

With the addition of 8 quads of 23.5° beams (shown in gray in Figure 7-3) the total number of beams available is 120, 96 of which are arranged in two cones (at 23.5° and 44.5°) to the hohlraum axis. These 96 beams are essentially half of the fully symmetric NIF and have a similar distribution of energy between inner (23.5°) and outer (44.5°)

cones of beams. With this configuration, independent pulse shaping on the inner and outer beams allows greater control of symmetry of the x-ray field seen by the imploding capsule, allowing higher convergence implosions more relevant to future ignition experiments. Experiments to tune time-dependent symmetry and time multiple shocks in imploding ignition capsules can make significant progress at this point. With the introduction of the 23.5° beams, direct-drive experiments such as shown in Figure 7-2c become more useful, the narrow cone angle of the 23.5° beams allowing the foil target to accelerate further before the finite focal spot size makes the experiment two dimensional. At this point there are 12 symmetric quads at the top and bottom of the chamber allowing experiments to be performed with vertical halfraums (Figure 7-2b) and up to 240 kJ incident through one LEH.

After 120 beams, the remaining beams are added in symmetric sets allowing increasing energy or backlighter flexibility in many experiments, until 192 beams are available with the full 16-fold rotational symmetry of the NIF.

The lower graphics in Figure 7-3 show the evolution of the baseline deployment sequence highlighting the absence of symmetry until completion of the 4th cluster. Quads of the same color in the lower part of the figure are from the same cluster and do not possess rotational symmetry about the vertical.

It must be reiterated that the sequence shown in Figure 7-3 is not cast in concrete, the flexibility inherent in the Mission-First strategy allows us to continue to optimize the sequence to meet mission priorities. The agreement on the deployment sequence is a milestone in the ICF management plan, which has formal change control procedures.